



US010990306B1

(12) **United States Patent**
Aggarwal et al.

(10) **Patent No.: US 10,990,306 B1**
(45) **Date of Patent: *Apr. 27, 2021**

(54) **BANDWIDTH SHARING FOR PAIRED STORAGE SYSTEMS**

(56) **References Cited**

U.S. PATENT DOCUMENTS

(71) Applicant: **PURE STORAGE, INC.**, Mountain View, CA (US)

5,706,210 A 1/1998 Kumano et al.
5,799,200 A 8/1998 Brant et al.
5,933,598 A 8/1999 Scales et al.

(72) Inventors: **Mudit Aggarwal**, Milpitas, CA (US);
Yuval Frandzel, Foster City, CA (US)

(Continued)

FOREIGN PATENT DOCUMENTS

(73) Assignee: **Pure Storage, Inc.**, Mountain View, CA (US)

EP 0725324 A2 8/1996
WO 2012087648 A2 6/2012

(Continued)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

OTHER PUBLICATIONS

This patent is subject to a terminal disclaimer.

Bellamy-McIntyre et al., "OpenID and the Enterprise: A Model-based Analysis of Single Sign-On Authentication", 15th IEEE International Enterprise Distributed Object Computing Conference (EDOC), Aug. 29, 2011, pp. 129-138, IEEE Computer Society, USA, DOI: 10.1109/EDOC.2011.26, ISBN: 978-1-4577-0362-1.

(Continued)

(21) Appl. No.: **16/855,175**

Primary Examiner — David E Martinez

(22) Filed: **Apr. 22, 2020**

(57) **ABSTRACT**

Related U.S. Application Data

(63) Continuation of application No. 16/171,453, filed on Oct. 26, 2018, now Pat. No. 10,671,302.

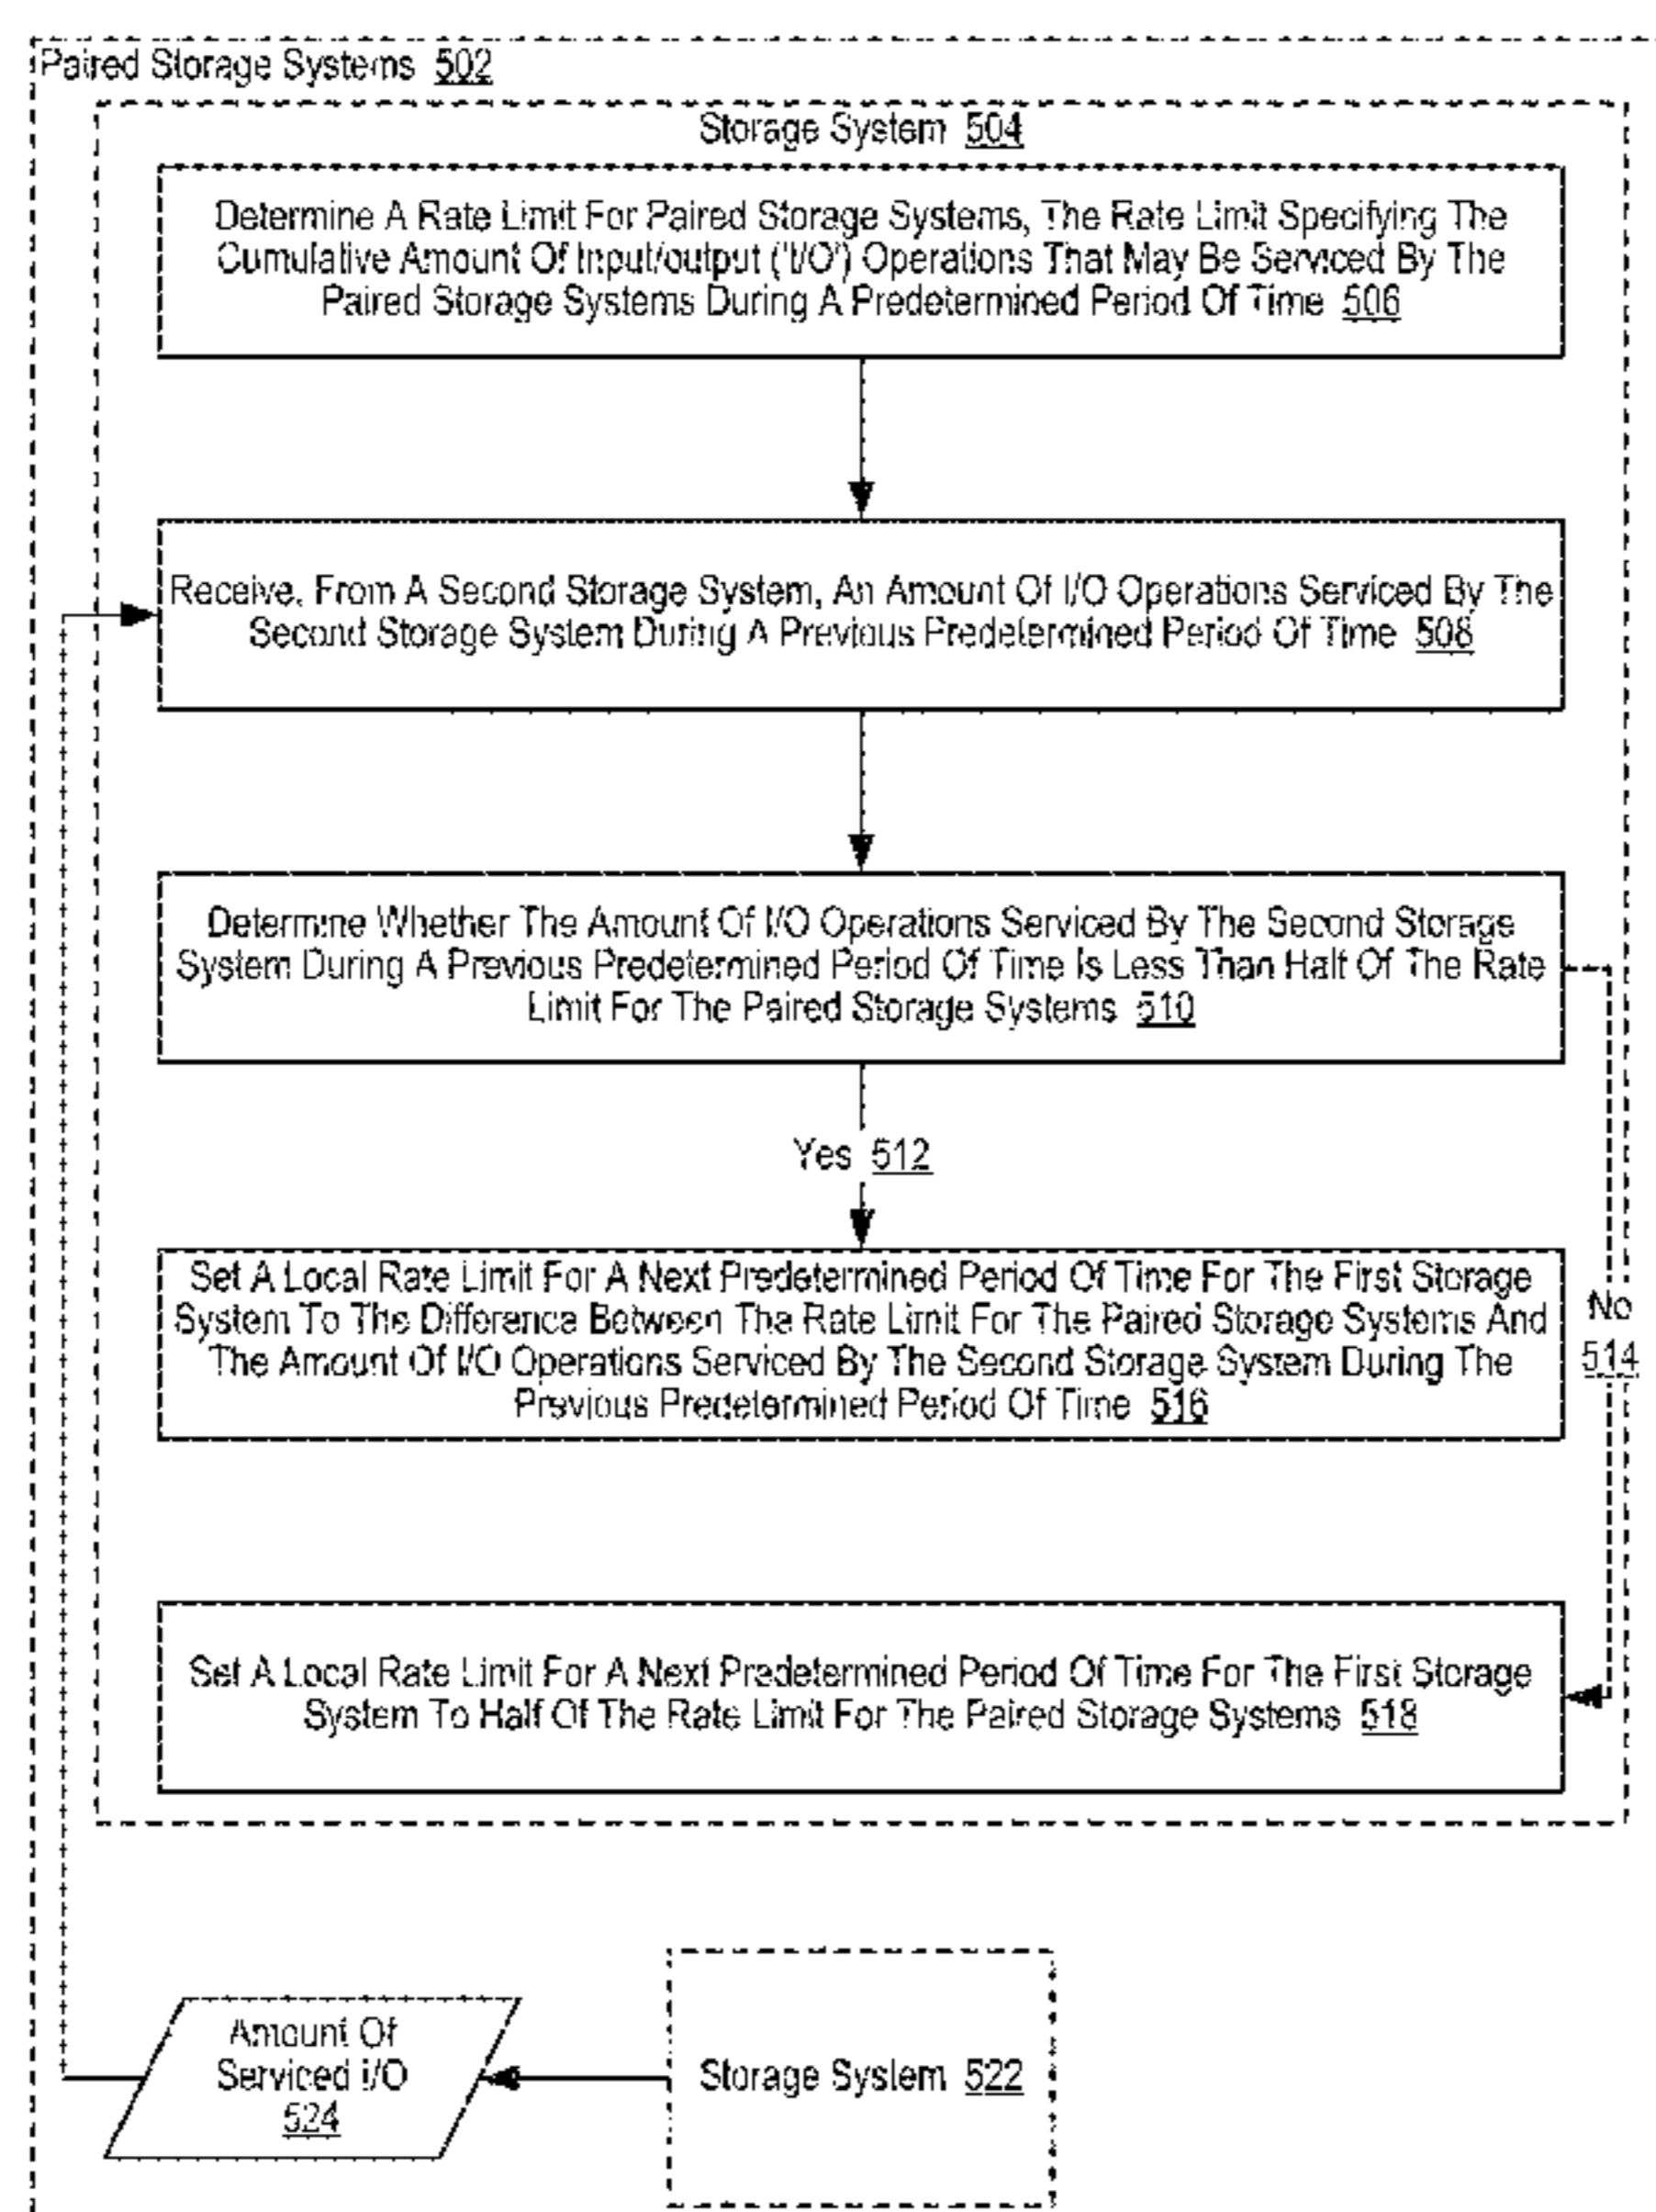
Applying a rate limit across a plurality of storage systems, including: determining a rate limit for paired storage systems; receiving, by a first storage system, an amount of I/O operations serviced by the second storage system during a previous predetermined period of time; determining whether the amount of I/O operations serviced by the second storage system is less than half of the rate limit for the paired storage systems; if so, setting local a rate limit for a next predetermined period of time for the first storage system to the difference between the rate limit for the paired storage systems and the amount of I/O operations serviced by the second storage system during the previous predetermined period of time; and otherwise, setting a local rate limit for a next predetermined period of time for the first storage system to half of the rate limit for the paired storage systems.

(51) **Int. Cl.**
G06F 3/06 (2006.01)
G06F 11/34 (2006.01)
G06F 13/16 (2006.01)

(52) **U.S. Cl.**
CPC **G06F 3/0632** (2013.01); **G06F 3/067** (2013.01); **G06F 3/0613** (2013.01); **G06F 3/0653** (2013.01)

(58) **Field of Classification Search**
CPC G06F 3/06; G06F 11/34; G06F 13/16
See application file for complete search history.

20 Claims, 17 Drawing Sheets



(56)

References Cited

U.S. PATENT DOCUMENTS

2008/0270678 A1 10/2008 Cornwell et al.
 2008/0282045 A1 11/2008 Biswas et al.
 2009/0037164 A1 2/2009 Gaither et al.
 2009/0077340 A1 3/2009 Johnson et al.
 2009/0100115 A1 4/2009 Park et al.
 2009/0198889 A1 8/2009 Ito et al.
 2010/0052625 A1 3/2010 Cagno et al.
 2010/0211723 A1 8/2010 Mukaida
 2010/0229166 A1 9/2010 Mopur et al.
 2010/0246266 A1 9/2010 Park et al.
 2010/0257142 A1 10/2010 Murphy et al.
 2010/0262764 A1 10/2010 Liu et al.
 2010/0262975 A1 10/2010 Reysa et al.
 2010/0325345 A1 12/2010 Ohno et al.
 2010/0332754 A1 12/2010 Lai et al.
 2011/0072290 A1 3/2011 Davis et al.
 2011/0125955 A1 5/2011 Chen
 2011/0131231 A1 6/2011 Haas et al.
 2011/0167221 A1 7/2011 Pangal et al.
 2012/0023144 A1 1/2012 Rub
 2012/0054264 A1 3/2012 Haugh et al.
 2012/0079318 A1 3/2012 Colgrove et al.
 2012/0131253 A1 5/2012 McKnight et al.
 2012/0131593 A1 5/2012 DePetro
 2012/0303919 A1 11/2012 Hu et al.
 2012/0311000 A1 12/2012 Post et al.
 2013/0007845 A1 1/2013 Chang et al.
 2013/0031414 A1 1/2013 Dhuse et al.
 2013/0036272 A1 2/2013 Nelson
 2013/0071087 A1 3/2013 Motiwala et al.
 2013/0145447 A1 6/2013 Maron
 2013/0191555 A1 7/2013 Liu
 2013/0198459 A1 8/2013 Joshi et al.
 2013/0205173 A1 8/2013 Yoneda
 2013/0219164 A1 8/2013 Hamid
 2013/0227201 A1 8/2013 Talagala et al.
 2013/0290607 A1 10/2013 Chang et al.
 2013/0311434 A1 11/2013 Jones
 2013/0318297 A1 11/2013 Jibbe et al.
 2013/0332614 A1 12/2013 Brunk et al.
 2014/0006358 A1 1/2014 Wang et al.
 2014/0020083 A1 1/2014 Fetik
 2014/0074850 A1 3/2014 Noel et al.
 2014/0082249 A1 3/2014 Pan et al.
 2014/0082715 A1 3/2014 Grajek et al.
 2014/0086146 A1 3/2014 Kim et al.
 2014/0090009 A1 3/2014 Li et al.
 2014/0096220 A1 4/2014 Pinto et al.
 2014/0101434 A1 4/2014 Senthurpandi et al.
 2014/0164774 A1 6/2014 Nord et al.
 2014/0173232 A1 6/2014 Reohr et al.
 2014/0195636 A1 7/2014 Karve et al.
 2014/0201512 A1 7/2014 Seethaler et al.
 2014/0201541 A1 7/2014 Paul et al.
 2014/0208155 A1 7/2014 Pan
 2014/0215590 A1 7/2014 Brand
 2014/0229654 A1 8/2014 Goss et al.
 2014/0230017 A1 8/2014 Saib
 2014/0244643 A1 8/2014 Basak et al.
 2014/0258526 A1 9/2014 Le Sant et al.
 2014/0282983 A1 9/2014 Ju et al.
 2014/0285917 A1 9/2014 Cudak et al.
 2014/0325262 A1 10/2014 Cooper et al.
 2014/0351627 A1 11/2014 Best et al.
 2014/0373104 A1 12/2014 Gaddam et al.
 2014/0373126 A1 12/2014 Hussain et al.
 2015/0026387 A1 1/2015 Sheredy et al.
 2015/0026582 A1 1/2015 Basak
 2015/0074463 A1 3/2015 Jacoby et al.
 2015/0089569 A1 3/2015 Sondhi et al.
 2015/0095515 A1 4/2015 Krithivas et al.
 2015/0113203 A1 4/2015 Dancho et al.
 2015/0121137 A1 4/2015 McKnight et al.
 2015/0134920 A1 5/2015 Anderson et al.
 2015/0149822 A1 5/2015 Coronado et al.

2015/0193169 A1 7/2015 Sundaram et al.
 2015/0378888 A1 12/2015 Zhang et al.
 2015/0379420 A1 12/2015 Basak et al.
 2016/0034372 A1 2/2016 Panda et al.
 2016/0065492 A1 3/2016 Hu et al.
 2016/0098323 A1 4/2016 Mutha et al.
 2016/0349992 A1 12/2016 Basak et al.
 2016/0350009 A1 12/2016 Cerreta et al.
 2016/0352720 A1 12/2016 Hu et al.
 2016/0352830 A1 12/2016 Borowiec et al.
 2016/0352834 A1 12/2016 Borowiec et al.
 2017/0052823 A1 2/2017 Hu et al.
 2017/0054799 A1 2/2017 Hu et al.
 2017/0083367 A1 3/2017 Chen et al.
 2017/0168729 A1 6/2017 Faulkner et al.
 2017/0364281 A1 12/2017 Chen et al.
 2018/0018339 A1 1/2018 Basak et al.
 2018/0048704 A1 2/2018 Li et al.
 2018/0103088 A1 4/2018 Blainey et al.
 2019/0056875 A1 2/2019 Nimura et al.

FOREIGN PATENT DOCUMENTS

WO 2013071087 A1 5/2013
 WO 2014110137 A1 7/2014
 WO 2016015008 A1 1/2016
 WO 2016190938 A1 12/2016
 WO 2016195759 A1 12/2016
 WO 2016195958 A1 12/2016
 WO 2016195961 A1 12/2016

OTHER PUBLICATIONS

ETSI, "Network Function Virtualisation (NFV); Resiliency Requirements", ETSI GS NFCV-REL 001, V1.1.1, Jan. 2015, 82 pages, etsi.org (online), URL: www.etsi.org/deliver/etsi_gs/NFV-REL/001_099/001/01.01.01_60/gs_NFV-REL001v010101p.pdf.
 Faith, "dictzip file format", GitHub.com (online), accessed Jul. 28, 2015, 1 page, URL: github.com/fidlej/idzip.
 Google Search of "storage array define" performed by the Examiner on Nov. 4, 2015 for U.S. Appl. No. 14/725,278, Results limited to entries dated before 2012, 1 page.
 Hota et al., "Capability-based Cryptographic Data Access Control in Cloud Computing", International Journal of Advanced Networking and Applications, col. 1, Issue 1, Aug. 2011, 10 pages, Eswar Publications, India.
 Hu et al., "Container Marking: Combining Data Placement, Garbage Collection and Wear Levelling for Flash", 19th Annual IEEE International Symposium on Modelling, Analysis, and Simulation of Computer and Telecommunications Systems, Jul. 25-27, 2011, 11 pages, ISBN: 978-0-7695-4430-4, DOI: 10.1109/MASCOTS.2011.50.
 International Search Report and Written Opinion, PCT/US2016/015006, dated Apr. 29, 2016, 12 pages.
 International Search Report and Written Opinion, PCT/US2016/015008, dated May 4, 2016, 12 pages.
 International Search Report and Written Opinion, PCT/US2016/016333, dated Jun. 8, 2016, 12 pages.
 International Search Report and Written Opinion, PCT/US2016/020410, dated Jul. 8, 2016, 12 pages.
 International Search Report and Written Opinion, PCT/US2016/032052, dated Aug. 30, 2016, 17 pages.
 International Search Report and Written Opinion, PCT/US2016/032084, dated Jul. 18, 2016, 12 pages.
 International Search Report and Written Opinion, PCT/US2016/035492, dated Aug. 17, 2016, 10 pages.
 International Search Report and Written Opinion, PCT/US2016/036693, dated Aug. 29, 2016, 10 pages.
 International Search Report and Written Opinion, PCT/US2016/038758, dated Oct. 7, 2016, 10 pages.
 International Search Report and Written Opinion, PCT/US2016/040393, dated Sep. 22, 2016, 10 pages.
 International Search Report and Written Opinion, PCT/US2016/044020, dated Sep. 30, 2016, 11 pages.

(56)

References Cited

OTHER PUBLICATIONS

International Search Report and Written Opinion, PCT/US2016/044874, dated Oct. 7, 2016, 11 pages.

International Search Report and Written Opinion, PCT/US2016/044875, dated Oct. 5, 2016, 13 pages.

International Search Report and Written Opinion, PCT/US2016/044876, dated Oct. 21, 2016, 12 pages.

International Search Report and Written Opinion, PCT/US2016/044877, dated Sep. 29, 2016, 13 pages.

Kong, "Using PCI Express As The Primary System Interconnect In Multiroot Compute, Storage, Communications And Embedded Systems", White Paper, IDT.com (online), Aug. 28, 2008, 12 pages, URL: www.idt.com/document/whp/idt-pcie-multi-root-white-paper.

Li et al., "Access Control for the Services Oriented Architecture", Proceedings of the 2007 ACM Workshop on Secure Web Services (SWS '07), Nov. 2007, pp. 9-17, ACM New York, NY.

Microsoft, "Hybrid for SharePoint Server 2013—Security Reference Architecture", Microsoft (online), Oct. 2014, 53 pages, URL: hybrid.office.com/img/Security_Reference_Architecture.pdf.

Microsoft, "Hybrid Identity Management", Microsoft (online), Apr. 2014, 2 pages, URL: [download.microsoft.com/download/E/A/E/](http://download.microsoft.com/download/E/A/E/EAE57CD1-A80B-423C-96BB-142FAAC630B9/Hybrid_Identity_Datasheet.pdf)

[EAE57CD1-A80B-423C-96BB-142FAAC630B9/Hybrid_Identity_Datasheet.pdf](http://download.microsoft.com/download/E/A/E/EAE57CD1-A80B-423C-96BB-142FAAC630B9/Hybrid_Identity_Datasheet.pdf).

Microsoft, "Hybrid Identity", Microsoft (online), Apr. 2014, 36 pages, URL: www.aka.ms/HybridIdentityWp.

PCMAG, "Storage Array Definition", Published May 10, 2013, URL: <http://web.archive.org/web/20130510121646/http://www.pcmag.com/encyclopedia/term/52091/storage-array>, 2 pages.

Storer et al., "Secure Data Deduplication", Proceedings of the 4th ACM International Workshop on Storage Security And Survivability (StorageSS'08), Oct. 2008, 10 pages, ACM New York, NY, USA, DOI: 10.1145/1456469.1456471.

Sweere, "Creating Storage Class Persistent Memory with NVDIMM", Published in Aug. 2013, Flash Memory Summit 2013, URL: http://www.flashmemorysummit.com/English/Collaterals/Proceedings/2013/20130814_T2_Sweere.pdf, 22 pages.

Techopedia, "What is a disk array", techopedia.com (online), Jan. 13, 2012, 1 page, URL: web.archive.org/web/20120113053358/http://www.techopedia.com/definition/1009/disk-array.

Webopedia, "What is a disk array", webopedia.com (online), May 26, 2011, 2 pages, URL: web.archive.org/web/20110526081214/http://www.webopedia.com/TERM/D/disk_array.html.

Wikipedia, "Convergent Encryption", Wikipedia.org (online), accessed Sep. 8, 2015, 2 pages, URL: en.wikipedia.org/wiki/Convergent_encryption.




FIG. 1A




FIG. 1B




FIG. 1C




FIG. 1D




FIG. 2A




FIG. 2B




FIG. 2C

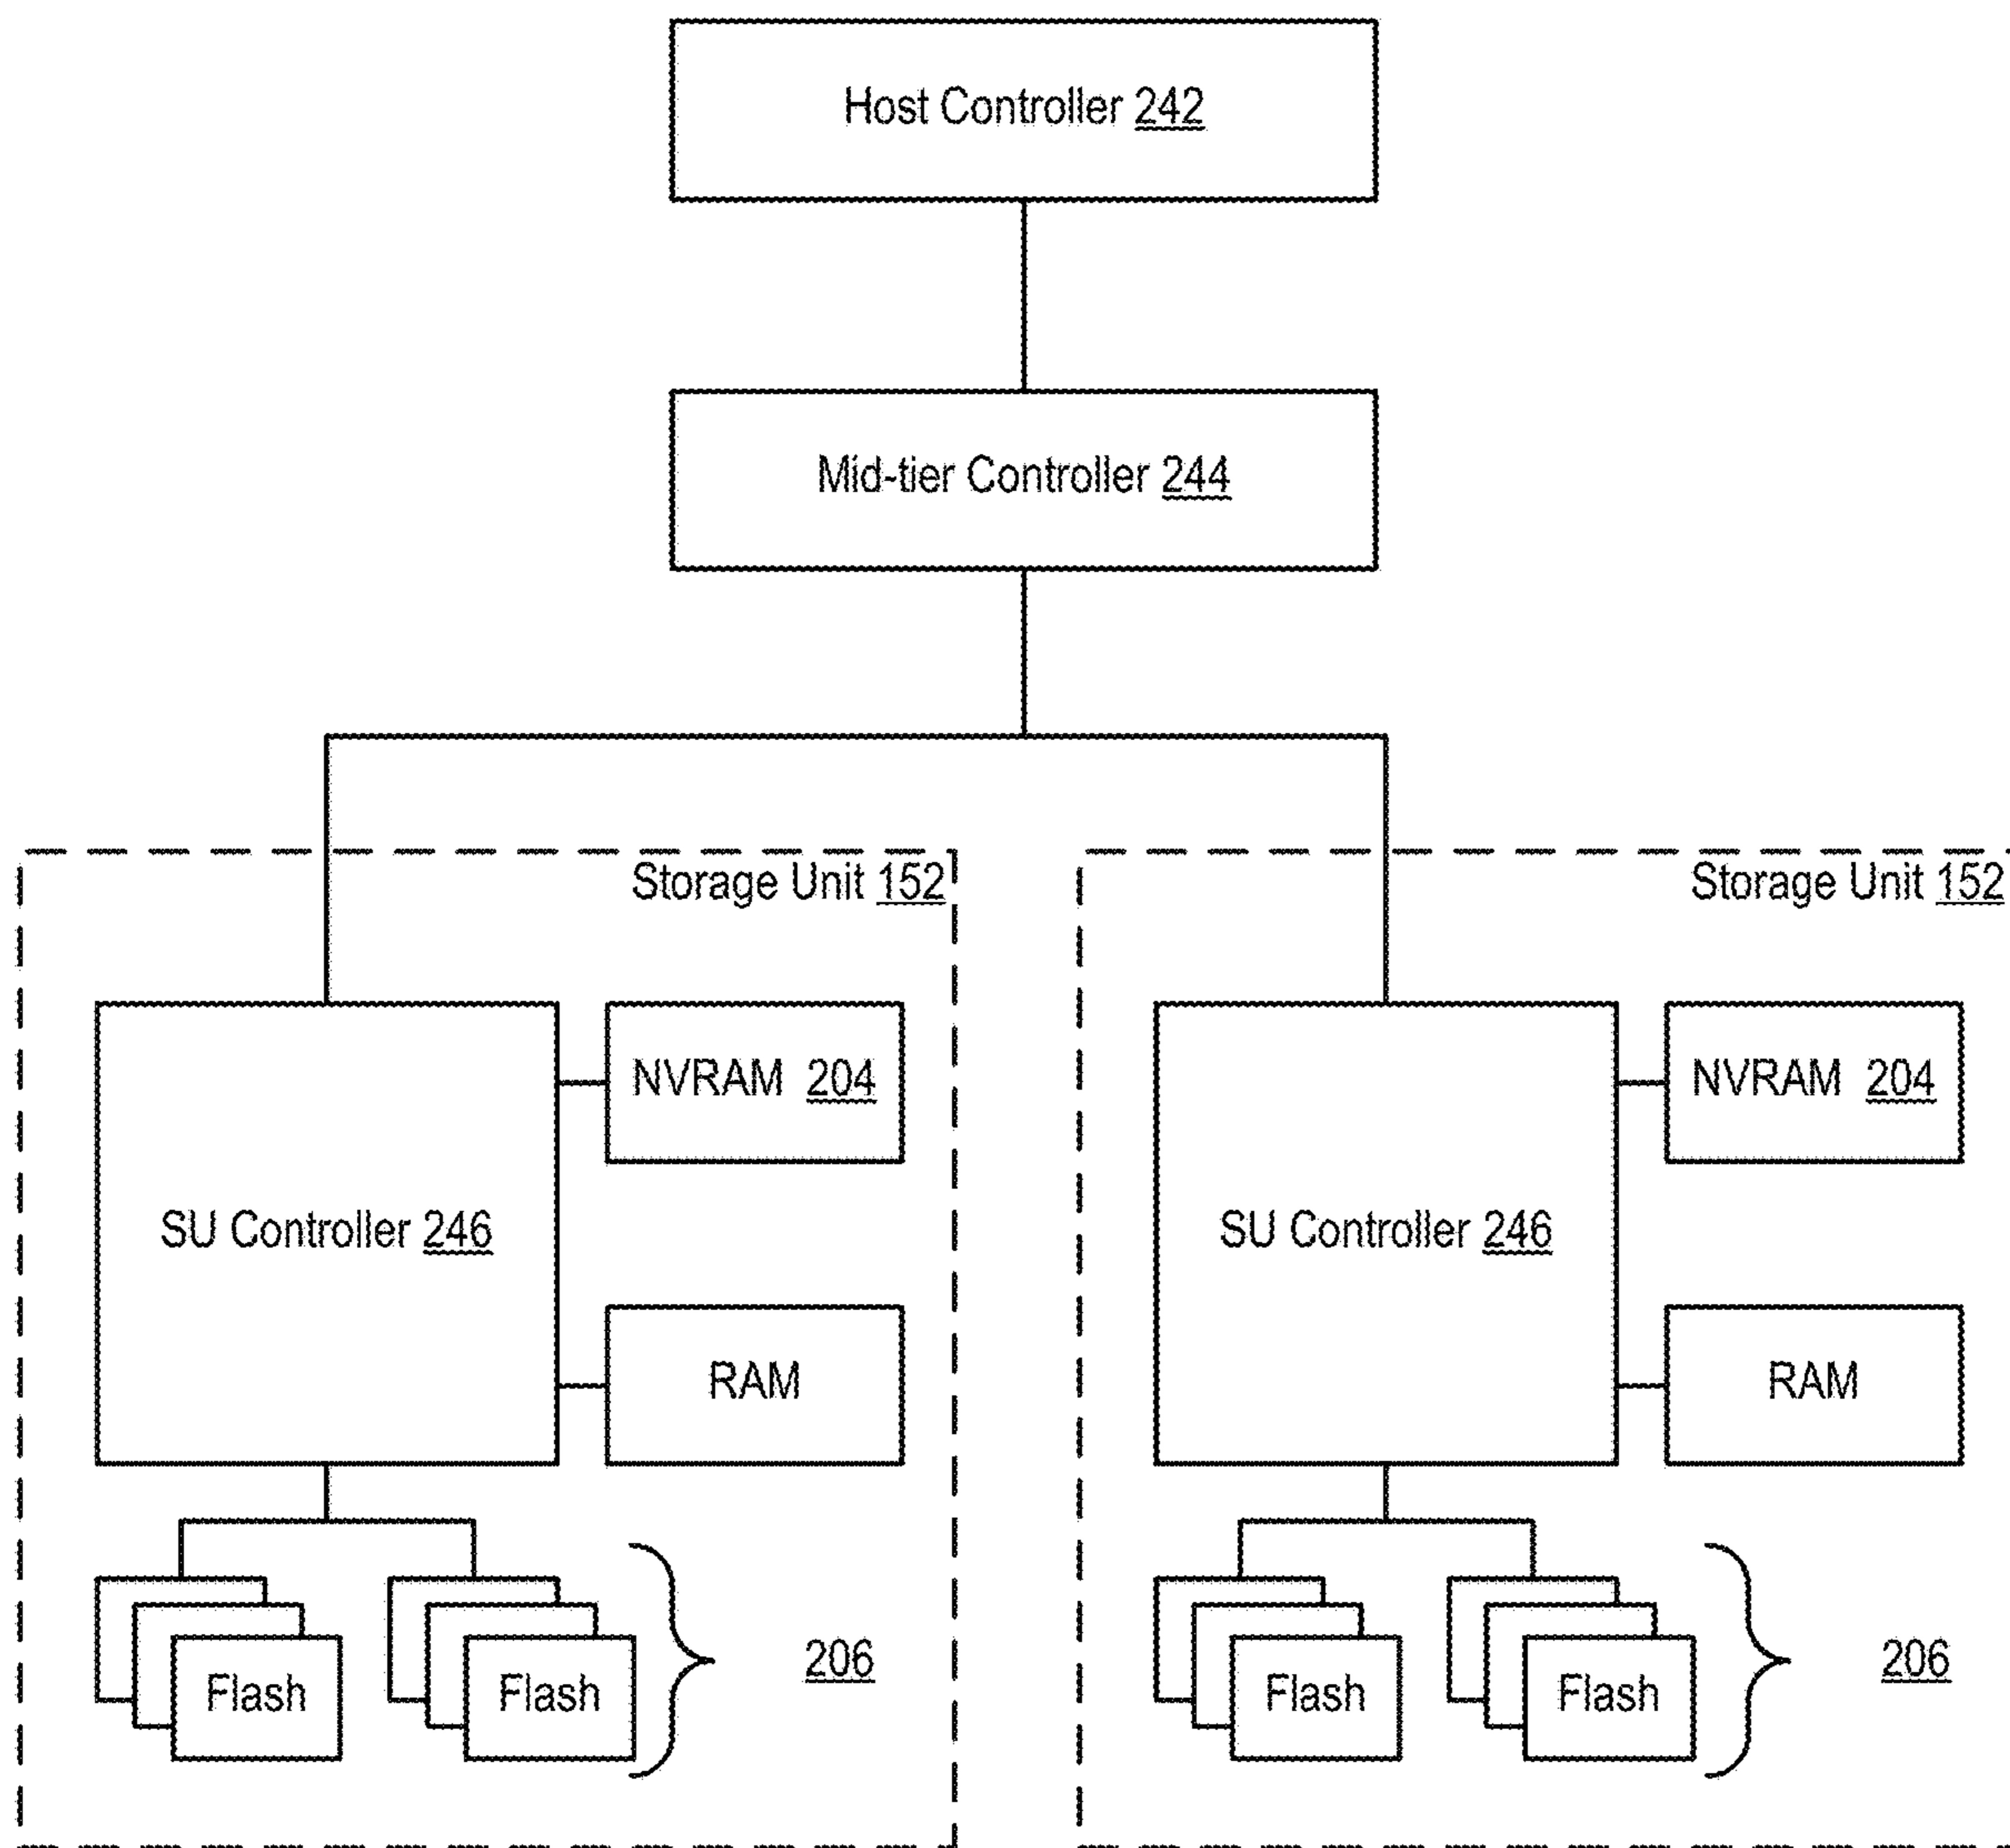


FIG. 2D




FIG. 2E

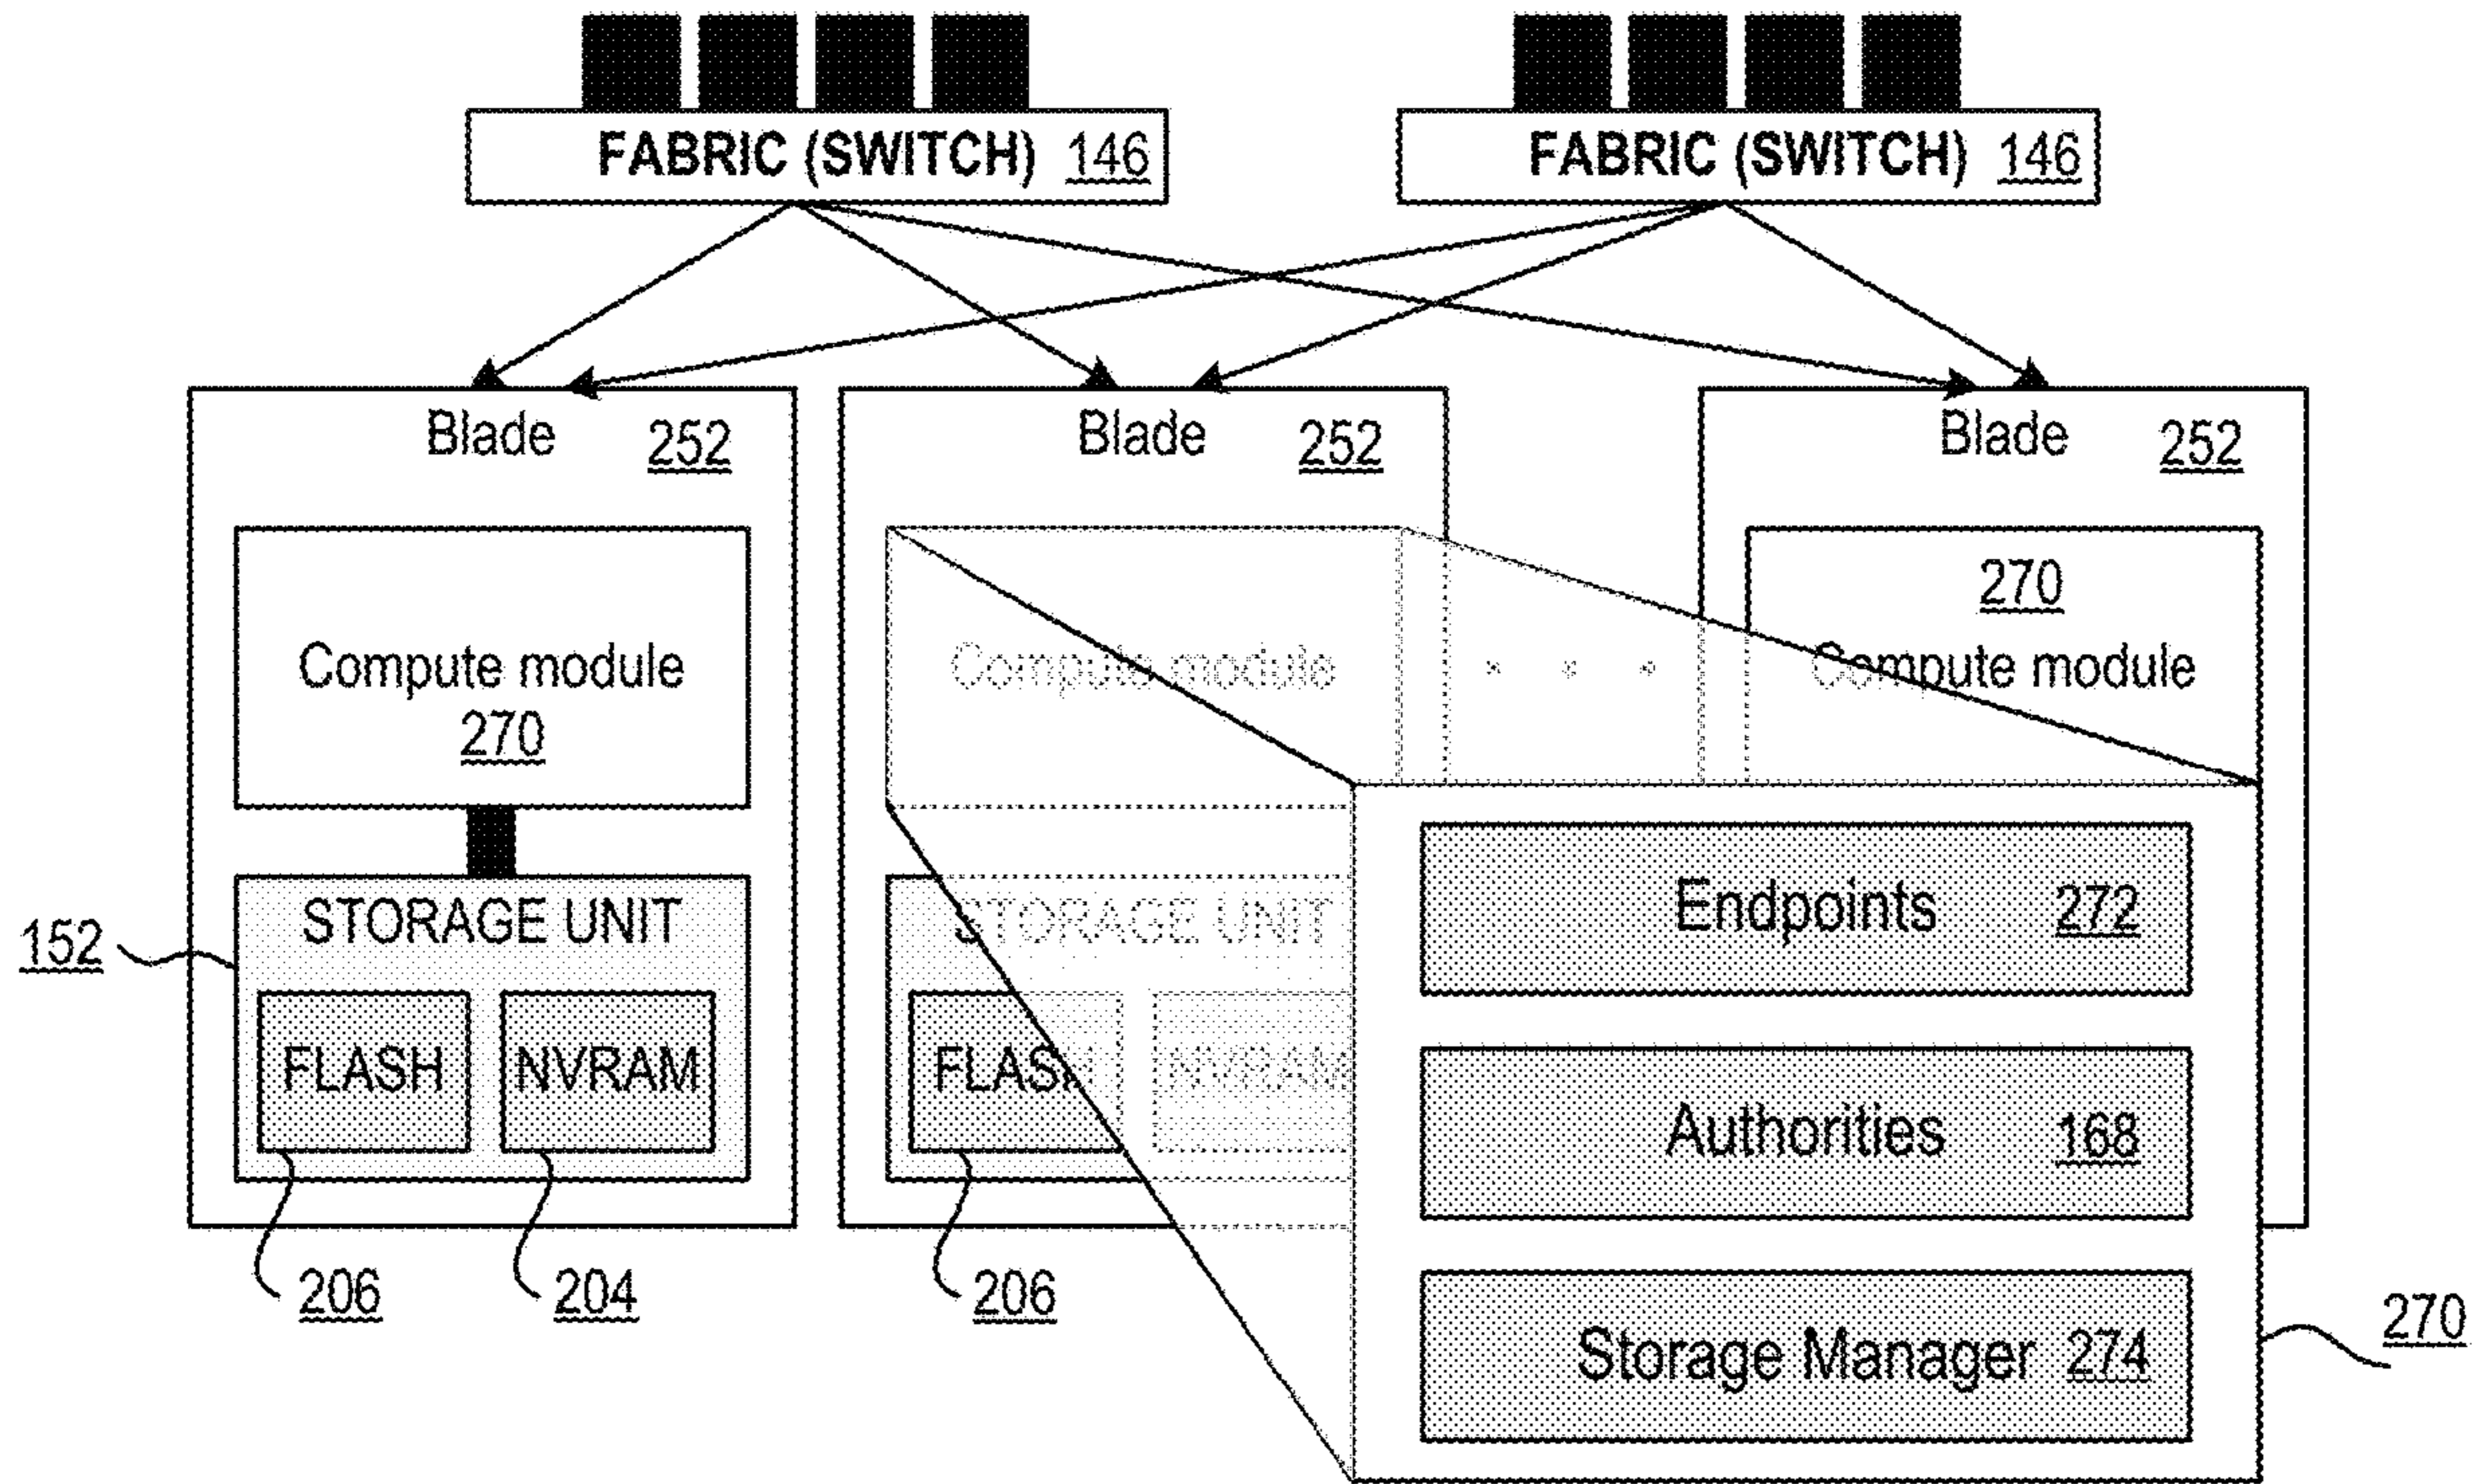


FIG. 2F




FIG. 2G




FIG. 3A




FIG. 3B




FIG. 4




FIG. 5




FIG. 6




FIG. 7




FIG. 8

BANDWIDTH SHARING FOR PAIRED STORAGE SYSTEMS

CROSS-REFERENCE TO RELATED APPLICATION

This application is a continuation application of and claims priority from U.S. patent application Ser. No. 16/171,453, filed Oct. 26, 2018.

BRIEF DESCRIPTION OF DRAWINGS

FIG. 1A illustrates a first example system for data storage in accordance with some implementations.

FIG. 1B illustrates a second example system for data storage in accordance with some implementations.

FIG. 1C illustrates a third example system for data storage in accordance with some implementations.

FIG. 1D illustrates a fourth example system for data storage in accordance with some implementations.

FIG. 2A is a perspective view of a storage cluster with multiple storage nodes and internal storage coupled to each storage node to provide network attached storage, in accordance with some embodiments.

FIG. 2B is a block diagram showing an interconnect switch coupling multiple storage nodes in accordance with some embodiments.

FIG. 2C is a multiple level block diagram, showing contents of a storage node and contents of one of the non-volatile solid state storage units in accordance with some embodiments.

FIG. 2D shows a storage server environment, which uses embodiments of the storage nodes and storage units of some previous figures in accordance with some embodiments.

FIG. 2E is a blade hardware block diagram, showing a control plane, compute and storage planes, and authorities interacting with underlying physical resources, in accordance with some embodiments.

FIG. 2F depicts elasticity software layers in blades of a storage cluster, in accordance with some embodiments.

FIG. 2G depicts authorities and storage resources in blades of a storage cluster, in accordance with some embodiments.

FIG. 3A sets forth a diagram of a storage system that is coupled for data communications with a cloud services provider in accordance with some embodiments of the present disclosure.

FIG. 3B sets forth a diagram of a storage system in accordance with some embodiments of the present disclosure.

FIG. 4 sets forth a block diagram illustrating a plurality of storage systems that support a pod according to some embodiments of the present disclosure.

For further explanation, FIG. 5 sets forth a flowchart illustrating an example method of applying a rate limit across a plurality of storage systems in accordance with some embodiments of the present disclosure.

For further explanation, FIG. 6 sets forth a flowchart illustrating an additional example method of applying a rate limit across a plurality of storage systems in accordance with some embodiments of the present disclosure.

For further explanation, FIG. 7 sets forth a flowchart illustrating an additional example method of applying a rate limit across a plurality of storage systems in accordance with some embodiments of the present disclosure.

For further explanation, FIG. 8 sets forth a flowchart illustrating an additional example method of applying a rate

limit across a plurality of storage systems in accordance with some embodiments of the present disclosure.

DESCRIPTION OF EMBODIMENTS

5

Example methods, apparatus, and products for applying a rate limit across a plurality of storage systems in accordance with embodiments of the present disclosure are described with reference to the accompanying drawings, beginning with FIG. 1A. FIG. 1A illustrates an example system for data storage, in accordance with some implementations. System **100** (also referred to as “storage system” herein) includes numerous elements for purposes of illustration rather than limitation. It may be noted that system **100** may include the same, more, or fewer elements configured in the same or different manner in other implementations.

System **100** includes a number of computing devices **164A-B**. Computing devices (also referred to as “client devices” herein) may be embodied, for example, a server in a data center, a workstation, a personal computer, a notebook, or the like. Computing devices **164A-B** may be coupled for data communications to one or more storage arrays **102A-B** through a storage area network (“SAN”) **158** or a local area network (“LAN”) **160**.

The SAN **158** may be implemented with a variety of data communications fabrics, devices, and protocols. For example, the fabrics for SAN **158** may include Fibre Channel, Ethernet, Infiniband, Serial Attached Small Computer System Interface (“SAS”), or the like. Data communications protocols for use with SAN **158** may include Advanced Technology Attachment (“ATA”), Fibre Channel Protocol, Small Computer System Interface (“SCSI”), Internet Small Computer System Interface (“iSCSI”), HyperSCSI, Non-Volatile Memory Express (“NVMe”) over Fabrics, or the like. It may be noted that SAN **158** is provided for illustration, rather than limitation. Other data communication couplings may be implemented between computing devices **164A-B** and storage arrays **102A-B**.

The LAN **160** may also be implemented with a variety of fabrics, devices, and protocols. For example, the fabrics for LAN **160** may include Ethernet (**802.3**), wireless (**802.11**), or the like. Data communication protocols for use in LAN **160** may include Transmission Control Protocol (“TCP”), User Datagram Protocol (“UDP”), Internet Protocol (“IP”), HyperText Transfer Protocol (“HTTP”), Wireless Access Protocol (“WAP”), Handheld Device Transport Protocol (“HDTP”), Session Initiation Protocol (“SIP”), Real Time Protocol (“RTP”), or the like.

Storage arrays **102A-B** may provide persistent data storage for the computing devices **164A-B**. Storage array **102A** may be contained in a chassis (not shown), and storage array **102B** may be contained in another chassis (not shown), in implementations. Storage array **102A** and **102B** may include one or more storage array controllers **110A-D** (also referred to as “controller” herein). A storage array controller **110A-D** may be embodied as a module of automated computing machinery comprising computer hardware, computer software, or a combination of computer hardware and software. In some implementations, the storage array controllers **110A-D** may be configured to carry out various storage tasks. Storage tasks may include writing data received from the computing devices **164A-B** to storage array **102A-B**, erasing data from storage array **102A-B**, retrieving data from storage array **102A-B** and providing data to computing devices **164A-B**, monitoring and reporting of disk utilization and performance, performing redundancy operations, such as Redundant Array of Independent Drives (“RAID”) or

RAID-like data redundancy operations, compressing data, encrypting data, and so forth.

Storage array controller **110A-D** may be implemented in a variety of ways, including as a Field Programmable Gate Array ('FPGA'), a Programmable Logic Chip ('PLC'), an Application Specific Integrated Circuit ('ASIC'), System-on-Chip ('SOC'), or any computing device that includes discrete components such as a processing device, central processing unit, computer memory, or various adapters. Storage array controller **110A-D** may include, for example, a data communications adapter configured to support communications via the SAN **158** or LAN **160**. In some implementations, storage array controller **110A-D** may be independently coupled to the LAN **160**. In implementations, storage array controller **110A-D** may include an I/O controller or the like that couples the storage array controller **110A-D** for data communications, through a midplane (not shown), to a persistent storage resource **170A-B** (also referred to as a "storage resource" herein). The persistent storage resource **170A-B** may include any number of storage drives **171A-F** (also referred to as "storage devices" herein) and any number of non-volatile Random Access Memory ('NVRAM') devices (not shown).

In some implementations, the NVRAM devices of a persistent storage resource **170A-B** may be configured to receive, from the storage array controller **110A-D**, data to be stored in the storage drives **171A-F**. In some examples, the data may originate from computing devices **164A-B**. In some examples, writing data to the NVRAM device may be carried out more quickly than directly writing data to the storage drive **171A-F**. In implementations, the storage array controller **110A-D** may be configured to utilize the NVRAM devices as a quickly accessible buffer for data destined to be written to the storage drives **171A-F**. Latency for write requests using NVRAM devices as a buffer may be improved relative to a system in which a storage array controller **110A-D** writes data directly to the storage drives **171A-F**. In some implementations, the NVRAM devices may be implemented with computer memory in the form of high bandwidth, low latency RAM. The NVRAM device is referred to as "non-volatile" because the NVRAM device may receive or include a unique power source that maintains the state of the RAM after main power loss to the NVRAM device. Such a power source may be a battery, one or more capacitors, or the like. In response to a power loss, the NVRAM device may be configured to write the contents of the RAM to a persistent storage, such as the storage drives **171A-F**.

In implementations, storage drive **171A-F** may refer to any device configured to record data persistently, where "persistently" or "persistent" refers as to a device's ability to maintain recorded data after loss of power. In some implementations, storage drive **171A-F** may correspond to non-disk storage media. For example, the storage drive **171A-F** may be one or more solid-state drives ('SSDs'), flash memory based storage, any type of solid-state non-volatile memory, or any other type of non-mechanical storage device. In other implementations, storage drive **171A-F** may include mechanical or spinning hard disk, such as hard-disk drives ('HDD').

In some implementations, the storage array controllers **110A-D** may be configured for offloading device management responsibilities from storage drive **171A-F** in storage array **102A-B**. For example, storage array controllers **110A-D** may manage control information that may describe the state of one or more memory blocks in the storage drives **171A-F**. The control information may indicate, for example,

that a particular memory block has failed and should no longer be written to, that a particular memory block contains boot code for a storage array controller **110A-D**, the number of program-erase ('P/E') cycles that have been performed on a particular memory block, the age of data stored in a particular memory block, the type of data that is stored in a particular memory block, and so forth. In some implementations, the control information may be stored with an associated memory block as metadata. In other implementations, the control information for the storage drives **171A-F** may be stored in one or more particular memory blocks of the storage drives **171A-F** that are selected by the storage array controller **110A-D**. The selected memory blocks may be tagged with an identifier indicating that the selected memory block contains control information. The identifier may be utilized by the storage array controllers **110A-D** in conjunction with storage drives **171A-F** to quickly identify the memory blocks that contain control information. For example, the storage controllers **110A-D** may issue a command to locate memory blocks that contain control information. It may be noted that control information may be so large that parts of the control information may be stored in multiple locations, that the control information may be stored in multiple locations for purposes of redundancy, for example, or that the control information may otherwise be distributed across multiple memory blocks in the storage drive **171A-F**.

In implementations, storage array controllers **110A-D** may offload device management responsibilities from storage drives **171A-F** of storage array **102A-B** by retrieving, from the storage drives **171A-F**, control information describing the state of one or more memory blocks in the storage drives **171A-F**. Retrieving the control information from the storage drives **171A-F** may be carried out, for example, by the storage array controller **110A-D** querying the storage drives **171A-F** for the location of control information for a particular storage drive **171A-F**. The storage drives **171A-F** may be configured to execute instructions that enable the storage drive **171A-F** to identify the location of the control information. The instructions may be executed by a controller (not shown) associated with or otherwise located on the storage drive **171A-F** and may cause the storage drive **171A-F** to scan a portion of each memory block to identify the memory blocks that store control information for the storage drives **171A-F**. The storage drives **171A-F** may respond by sending a response message to the storage array controller **110A-D** that includes the location of control information for the storage drive **171A-F**. Responsive to receiving the response message, storage array controllers **110A-D** may issue a request to read data stored at the address associated with the location of control information for the storage drives **171A-F**.

In other implementations, the storage array controllers **110A-D** may further offload device management responsibilities from storage drives **171A-F** by performing, in response to receiving the control information, a storage drive management operation. A storage drive management operation may include, for example, an operation that is typically performed by the storage drive **171A-F** (e.g., the controller (not shown) associated with a particular storage drive **171A-F**). A storage drive management operation may include, for example, ensuring that data is not written to failed memory blocks within the storage drive **171A-F**, ensuring that data is written to memory blocks within the storage drive **171A-F** in such a way that adequate wear leveling is achieved, and so forth.

In implementations, storage array 102A-B may implement two or more storage array controllers 110A-D. For example, storage array 102A may include storage array controllers 110A and storage array controllers 110B. At a given instance, a single storage array controller 110A-D (e.g., storage array controller 110A) of a storage system 100 may be designated with primary status (also referred to as “primary controller” herein), and other storage array controllers 110A-D (e.g., storage array controller 110A) may be designated with secondary status (also referred to as “secondary controller” herein). The primary controller may have particular rights, such as permission to alter data in persistent storage resource 170A-B (e.g., writing data to persistent storage resource 170A-B). At least some of the rights of the primary controller may supersede the rights of the secondary controller. For instance, the secondary controller may not have permission to alter data in persistent storage resource 170A-B when the primary controller has the right. The status of storage array controllers 110A-D may change. For example, storage array controller 110A may be designated with secondary status, and storage array controller 110B may be designated with primary status.

In some implementations, a primary controller, such as storage array controller 110A, may serve as the primary controller for one or more storage arrays 102A-B, and a second controller, such as storage array controller 110B, may serve as the secondary controller for the one or more storage arrays 102A-B. For example, storage array controller 110A may be the primary controller for storage array 102A and storage array 102B, and storage array controller 110B may be the secondary controller for storage array 102A and 102B. In some implementations, storage array controllers 110C and 110D (also referred to as “storage processing modules”) may neither have primary or secondary status. Storage array controllers 110C and 110D, implemented as storage processing modules, may act as a communication interface between the primary and secondary controllers (e.g., storage array controllers 110A and 110B, respectively) and storage array 102B. For example, storage array controller 110A of storage array 102A may send a write request, via SAN 158, to storage array 102B. The write request may be received by both storage array controllers 110C and 110D of storage array 102B. Storage array controllers 110C and 110D facilitate the communication, e.g., send the write request to the appropriate storage drive 171A-F. It may be noted that in some implementations storage processing modules may be used to increase the number of storage drives controlled by the primary and secondary controllers.

In implementations, storage array controllers 110A-D are communicatively coupled, via a midplane (not shown), to one or more storage drives 171A-F and to one or more NVRAM devices (not shown) that are included as part of a storage array 102A-B. The storage array controllers 110A-D may be coupled to the midplane via one or more data communication links and the midplane may be coupled to the storage drives 171A-F and the NVRAM devices via one or more data communications links. The data communications links described herein are collectively illustrated by data communications links 108A-D and may include a Peripheral Component Interconnect Express (‘PCIe’) bus, for example.

FIG. 1B illustrates an example system for data storage, in accordance with some implementations. Storage array controller 101 illustrated in FIG. 1B may similar to the storage array controllers 110A-D described with respect to FIG. 1A. In one example, storage array controller 101 may be similar

to storage array controller 110A or storage array controller 110B. Storage array controller 101 includes numerous elements for purposes of illustration rather than limitation. It may be noted that storage array controller 101 may include the same, more, or fewer elements configured in the same or different manner in other implementations. It may be noted that elements of FIG. 1A may be included below to help illustrate features of storage array controller 101.

Storage array controller 101 may include one or more processing devices 104 and random access memory (‘RAM’) 111. Processing device 104 (or controller 101) represents one or more general-purpose processing devices such as a microprocessor, central processing unit, or the like. More particularly, the processing device 104 (or controller 101) may be a complex instruction set computing (‘CISC’) microprocessor, reduced instruction set computing (‘RISC’) microprocessor, very long instruction word (‘VLIW’) microprocessor, or a processor implementing other instruction sets or processors implementing a combination of instruction sets. The processing device 104 (or controller 101) may also be one or more special-purpose processing devices such as an application specific integrated circuit (‘ASIC’), a field programmable gate array (‘FPGA’), a digital signal processor (‘DSP’), network processor, or the like.

The processing device 104 may be connected to the RAM 111 via a data communications link 106, which may be embodied as a high speed memory bus such as a Double-Data Rate 4 (‘DDR4’) bus. Stored in RAM 111 is an operating system 112. In some implementations, instructions 113 are stored in RAM 111. Instructions 113 may include computer program instructions for performing operations in a direct-mapped flash storage system. In one embodiment, a direct-mapped flash storage system is one that addresses data blocks within flash drives directly and without an address translation performed by the storage controllers of the flash drives.

In implementations, storage array controller 101 includes one or more host bus adapters 103A-C that are coupled to the processing device 104 via a data communications link 105A-C. In implementations, host bus adapters 103A-C may be computer hardware that connects a host system (e.g., the storage array controller) to other network and storage arrays. In some examples, host bus adapters 103A-C may be a Fibre Channel adapter that enables the storage array controller 101 to connect to a SAN, an Ethernet adapter that enables the storage array controller 101 to connect to a LAN, or the like. Host bus adapters 103A-C may be coupled to the processing device 104 via a data communications link 105A-C such as, for example, a PCIe bus.

In implementations, storage array controller 101 may include a host bus adapter 114 that is coupled to an expander 115. The expander 115 may be used to attach a host system to a larger number of storage drives. The expander 115 may, for example, be a SAS expander utilized to enable the host bus adapter 114 to attach to storage drives in an implementation where the host bus adapter 114 is embodied as a SAS controller.

In implementations, storage array controller 101 may include a switch 116 coupled to the processing device 104 via a data communications link 109. The switch 116 may be a computer hardware device that can create multiple endpoints out of a single endpoint, thereby enabling multiple devices to share a single endpoint. The switch 116 may, for example, be a PCIe switch that is coupled to a PCIe bus (e.g., data communications link 109) and presents multiple PCIe connection points to the midplane.

In implementations, storage array controller **101** includes a data communications link **107** for coupling the storage array controller **101** to other storage array controllers. In some examples, data communications link **107** may be a QuickPath Interconnect (QPI) interconnect.

A traditional storage system that uses traditional flash drives may implement a process across the flash drives that are part of the traditional storage system. For example, a higher level process of the storage system may initiate and control a process across the flash drives. However, a flash drive of the traditional storage system may include its own storage controller that also performs the process. Thus, for the traditional storage system, a higher level process (e.g., initiated by the storage system) and a lower level process (e.g., initiated by a storage controller of the storage system) may both be performed.

To resolve various deficiencies of a traditional storage system, operations may be performed by higher level processes and not by the lower level processes. For example, the flash storage system may include flash drives that do not include storage controllers that provide the process. Thus, the operating system of the flash storage system itself may initiate and control the process. This may be accomplished by a direct-mapped flash storage system that addresses data blocks within the flash drives directly and without an address translation performed by the storage controllers of the flash drives.

The operating system of the flash storage system may identify and maintain a list of allocation units across multiple flash drives of the flash storage system. The allocation units may be entire erase blocks or multiple erase blocks. The operating system may maintain a map or address range that directly maps addresses to erase blocks of the flash drives of the flash storage system.

Direct mapping to the erase blocks of the flash drives may be used to rewrite data and erase data. For example, the operations may be performed on one or more allocation units that include a first data and a second data where the first data is to be retained and the second data is no longer being used by the flash storage system. The operating system may initiate the process to write the first data to new locations within other allocation units and erasing the second data and marking the allocation units as being available for use for subsequent data. Thus, the process may only be performed by the higher level operating system of the flash storage system without an additional lower level process being performed by controllers of the flash drives.

Advantages of the process being performed only by the operating system of the flash storage system include increased reliability of the flash drives of the flash storage system as unnecessary or redundant write operations are not being performed during the process. One possible point of novelty here is the concept of initiating and controlling the process at the operating system of the flash storage system. In addition, the process can be controlled by the operating system across multiple flash drives. This is contrast to the process being performed by a storage controller of a flash drive.

A storage system can consist of two storage array controllers that share a set of drives for failover purposes, or it could consist of a single storage array controller that provides a storage service that utilizes multiple drives, or it could consist of a distributed network of storage array controllers each with some number of drives or some amount of Flash storage where the storage array controllers in the network collaborate to provide a complete storage

service and collaborate on various aspects of a storage service including storage allocation and garbage collection.

FIG. **1C** illustrates a third example system **117** for data storage in accordance with some implementations. System **117** (also referred to as “storage system” herein) includes numerous elements for purposes of illustration rather than limitation. It may be noted that system **117** may include the same, more, or fewer elements configured in the same or different manner in other implementations.

In one embodiment, system **117** includes a dual Peripheral Component Interconnect (‘PCI’) flash storage device **118** with separately addressable fast write storage. System **117** may include a storage controller **119**. In one embodiment, storage controller **119A-D** may be a CPU, ASIC, FPGA, or any other circuitry that may implement control structures necessary according to the present disclosure. In one embodiment, system **117** includes flash memory devices (e.g., including flash memory devices **120a-n**), operatively coupled to various channels of the storage device controller **119**. Flash memory devices **120a-n**, may be presented to the controller **119A-D** as an addressable collection of Flash pages, erase blocks, and/or control elements sufficient to allow the storage device controller **119A-D** to program and retrieve various aspects of the Flash. In one embodiment, storage device controller **119A-D** may perform operations on flash memory devices **120a-n** including storing and retrieving data content of pages, arranging and erasing any blocks, tracking statistics related to the use and reuse of Flash memory pages, erase blocks, and cells, tracking and predicting error codes and faults within the Flash memory, controlling voltage levels associated with programming and retrieving contents of Flash cells, etc.

In one embodiment, system **117** may include RAM **121** to store separately addressable fast-write data. In one embodiment, RAM **121** may be one or more separate discrete devices. In another embodiment, RAM **121** may be integrated into storage device controller **119A-D** or multiple storage device controllers. The RAM **121** may be utilized for other purposes as well, such as temporary program memory for a processing device (e.g., a CPU) in the storage device controller **119**.

In one embodiment, system **117** may include a stored energy device **122**, such as a rechargeable battery or a capacitor. Stored energy device **122** may store energy sufficient to power the storage device controller **119**, some amount of the RAM (e.g., RAM **121**), and some amount of Flash memory (e.g., Flash memory **120a-120n**) for sufficient time to write the contents of RAM to Flash memory. In one embodiment, storage device controller **119A-D** may write the contents of RAM to Flash Memory if the storage device controller detects loss of external power.

In one embodiment, system **117** includes two data communications links **123a**, **123b**. In one embodiment, data communications links **123a**, **123b** may be PCI interfaces. In another embodiment, data communications links **123a**, **123b** may be based on other communications standards (e.g., HyperTransport, InfiniBand, etc.). Data communications links **123a**, **123b** may be based on non-volatile memory express (‘NVMe’) or NVMe over fabrics (‘NVMe-f’) specifications that allow external connection to the storage device controller **119A-D** from other components in the storage system **117**. It should be noted that data communications links may be interchangeably referred to herein as PCI buses for convenience.

System **117** may also include an external power source (not shown), which may be provided over one or both data communications links **123a**, **123b**, or which may be pro-

vided separately. An alternative embodiment includes a separate Flash memory (not shown) dedicated for use in storing the content of RAM **121**. The storage device controller **119A-D** may present a logical device over a PCI bus which may include an addressable fast-write logical device, or a distinct part of the logical address space of the storage device **118**, which may be presented as PCI memory or as persistent storage. In one embodiment, operations to store into the device are directed into the RAM **121**. On power failure, the storage device controller **119A-D** may write stored content associated with the addressable fast-write logical storage to Flash memory (e.g., Flash memory **120a-n**) for long-term persistent storage.

In one embodiment, the logical device may include some presentation of some or all of the content of the Flash memory devices **120a-n**, where that presentation allows a storage system including a storage device **118** (e.g., storage system **117**) to directly address Flash memory pages and directly reprogram erase blocks from storage system components that are external to the storage device through the PCI bus. The presentation may also allow one or more of the external components to control and retrieve other aspects of the Flash memory including some or all of: tracking statistics related to use and reuse of Flash memory pages, erase blocks, and cells across all the Flash memory devices; tracking and predicting error codes and faults within and across the Flash memory devices; controlling voltage levels associated with programming and retrieving contents of Flash cells; etc.

In one embodiment, the stored energy device **122** may be sufficient to ensure completion of in-progress operations to the Flash memory devices **120a-120n** stored energy device **122** may power storage device controller **119A-D** and associated Flash memory devices (e.g., **120a-n**) for those operations, as well as for the storing of fast-write RAM to Flash memory. Stored energy device **122** may be used to store accumulated statistics and other parameters kept and tracked by the Flash memory devices **120a-n** and/or the storage device controller **119**. Separate capacitors or stored energy devices (such as smaller capacitors near or embedded within the Flash memory devices themselves) may be used for some or all of the operations described herein.

Various schemes may be used to track and optimize the life span of the stored energy component, such as adjusting voltage levels over time, partially discharging the storage energy device **122** to measure corresponding discharge characteristics, etc. If the available energy decreases over time, the effective available capacity of the addressable fast-write storage may be decreased to ensure that it can be written safely based on the currently available stored energy.

FIG. **1D** illustrates a third example system **124** for data storage in accordance with some implementations. In one embodiment, system **124** includes storage controllers **125a**, **125b**. In one embodiment, storage controllers **125a**, **125b** are operatively coupled to Dual PCI storage devices **119a**, **119b** and **119c**, **119d**, respectively. Storage controllers **125a**, **125b** may be operatively coupled (e.g., via a storage network **130**) to some number of host computers **127a-n**.

In one embodiment, two storage controllers (e.g., **125a** and **125b**) provide storage services, such as a SCS) block storage array, a file server, an object server, a database, or data analytics service, etc. The storage controllers **125a**, **125b** may provide services through some number of network interfaces (e.g., **126a-d**) to host computers **127a-n** outside of the storage system **124**. Storage controllers **125a**, **125b** may provide integrated services or an application entirely within the storage system **124**, forming a converged storage and

compute system. The storage controllers **125a**, **125b** may utilize the fast write memory within or across storage devices **119a-d** to journal in progress operations to ensure the operations are not lost on a power failure, storage controller removal, storage controller or storage system shutdown, or some fault of one or more software or hardware components within the storage system **124**.

In one embodiment, controllers **125a**, **125b** operate as PCI masters to one or the other PCI buses **128a**, **128b**. In another embodiment, **128a** and **128b** may be based on other communications standards (e.g., HyperTransport, InfiniBand, etc.). Other storage system embodiments may operate storage controllers **125a**, **125b** as multi-masters for both PCI buses **128a**, **128b**. Alternately, a PCI/NVMe/NVMf switching infrastructure or fabric may connect multiple storage controllers. Some storage system embodiments may allow storage devices to communicate with each other directly rather than communicating only with storage controllers. In one embodiment, a storage device controller **119a** may be operable under direction from a storage controller **125a** to synthesize and transfer data to be stored into Flash memory devices from data that has been stored in RAM (e.g., RAM **121** of FIG. **1C**). For example, a recalculated version of RAM content may be transferred after a storage controller has determined that an operation has fully committed across the storage system, or when fast-write memory on the device has reached a certain used capacity, or after a certain amount of time, to ensure improve safety of the data or to release addressable fast-write capacity for reuse. This mechanism may be used, for example, to avoid a second transfer over a bus (e.g., **128a**, **128b**) from the storage controllers **125a**, **125b**. In one embodiment, a recalculation may include compressing data, attaching indexing or other metadata, combining multiple data segments together, performing erasure code calculations, etc.

In one embodiment, under direction from a storage controller **125a**, **125b**, a storage device controller **119a**, **119b** may be operable to calculate and transfer data to other storage devices from data stored in RAM (e.g., RAM **121** of FIG. **1C**) without involvement of the storage controllers **125a**, **125b**. This operation may be used to mirror data stored in one controller **125a** to another controller **125b**, or it could be used to offload compression, data aggregation, and/or erasure coding calculations and transfers to storage devices to reduce load on storage controllers or the storage controller interface **129a**, **129b** to the PCI bus **128a**, **128b**.

A storage device controller **119A-D** may include mechanisms for implementing high availability primitives for use by other parts of a storage system external to the Dual PCI storage device **118**. For example, reservation or exclusion primitives may be provided so that, in a storage system with two storage controllers providing a highly available storage service, one storage controller may prevent the other storage controller from accessing or continuing to access the storage device. This could be used, for example, in cases where one controller detects that the other controller is not functioning properly or where the interconnect between the two storage controllers may itself not be functioning properly.

In one embodiment, a storage system for use with Dual PCI direct mapped storage devices with separately addressable fast write storage includes systems that manage erase blocks or groups of erase blocks as allocation units for storing data on behalf of the storage service, or for storing metadata (e.g., indexes, logs, etc.) associated with the storage service, or for proper management of the storage system itself. Flash pages, which may be a few kilobytes in size, may be written as data arrives or as the storage system is to

persist data for long intervals of time (e.g., above a defined threshold of time). To commit data more quickly, or to reduce the number of writes to the Flash memory devices, the storage controllers may first write data into the separately addressable fast write storage on one more storage devices.

In one embodiment, the storage controllers **125a**, **125b** may initiate the use of erase blocks within and across storage devices (e.g., **118**) in accordance with an age and expected remaining lifespan of the storage devices, or based on other statistics. The storage controllers **125a**, **125b** may initiate garbage collection and data migration data between storage devices in accordance with pages that are no longer needed as well as to manage Flash page and erase block lifespans and to manage overall system performance.

In one embodiment, the storage system **124** may utilize mirroring and/or erasure coding schemes as part of storing data into addressable fast write storage and/or as part of writing data into allocation units associated with erase blocks. Erasure codes may be used across storage devices, as well as within erase blocks or allocation units, or within and across Flash memory devices on a single storage device, to provide redundancy against single or multiple storage device failures or to protect against internal corruptions of Flash memory pages resulting from Flash memory operations or from degradation of Flash memory cells. Mirroring and erasure coding at various levels may be used to recover from multiple types of failures that occur separately or in combination.

The embodiments depicted with reference to FIGS. **2A-G** illustrate a storage cluster that stores user data, such as user data originating from one or more user or client systems or other sources external to the storage cluster. The storage cluster distributes user data across storage nodes housed within a chassis, or across multiple chassis, using erasure coding and redundant copies of metadata. Erasure coding refers to a method of data protection or reconstruction in which data is stored across a set of different locations, such as disks, storage nodes or geographic locations. Flash memory is one type of solid-state memory that may be integrated with the embodiments, although the embodiments may be extended to other types of solid-state memory or other storage medium, including non-solid state memory. Control of storage locations and workloads are distributed across the storage locations in a clustered peer-to-peer system. Tasks such as mediating communications between the various storage nodes, detecting when a storage node has become unavailable, and balancing I/Os (inputs and outputs) across the various storage nodes, are all handled on a distributed basis. Data is laid out or distributed across multiple storage nodes in data fragments or stripes that support data recovery in some embodiments. Ownership of data can be reassigned within a cluster, independent of input and output patterns. This architecture described in more detail below allows a storage node in the cluster to fail, with the system remaining operational, since the data can be reconstructed from other storage nodes and thus remain available for input and output operations. In various embodiments, a storage node may be referred to as a cluster node, a blade, or a server.

The storage cluster may be contained within a chassis, i.e., an enclosure housing one or more storage nodes. A mechanism to provide power to each storage node, such as a power distribution bus, and a communication mechanism, such as a communication bus that enables communication between the storage nodes are included within the chassis. The storage cluster can run as an independent system in one

location according to some embodiments. In one embodiment, a chassis contains at least two instances of both the power distribution and the communication bus which may be enabled or disabled independently. The internal communication bus may be an Ethernet bus, however, other technologies such as PCIe, InfiniBand, and others, are equally suitable. The chassis provides a port for an external communication bus for enabling communication between multiple chassis, directly or through a switch, and with client systems. The external communication may use a technology such as Ethernet, InfiniBand, Fibre Channel, etc. In some embodiments, the external communication bus uses different communication bus technologies for inter-chassis and client communication. If a switch is deployed within or between chassis, the switch may act as a translation between multiple protocols or technologies. When multiple chassis are connected to define a storage cluster, the storage cluster may be accessed by a client using either proprietary interfaces or standard interfaces such as network file system ('NFS'), common internet file system ('CIFS'), small computer system interface ('SCSI') or hypertext transfer protocol ('HTTP'). Translation from the client protocol may occur at the switch, chassis external communication bus or within each storage node. In some embodiments, multiple chassis may be coupled or connected to each other through an aggregator switch. A portion and/or all of the coupled or connected chassis may be designated as a storage cluster. As discussed above, each chassis can have multiple blades, each blade has a media access control ('MAC') address, but the storage cluster is presented to an external network as having a single cluster IP address and a single MAC address in some embodiments.

Each storage node may be one or more storage servers and each storage server is connected to one or more non-volatile solid state memory units, which may be referred to as storage units or storage devices. One embodiment includes a single storage server in each storage node and between one to eight non-volatile solid state memory units, however this one example is not meant to be limiting. The storage server may include a processor, DRAM and interfaces for the internal communication bus and power distribution for each of the power buses. Inside the storage node, the interfaces and storage unit share a communication bus, e.g., PCI Express, in some embodiments. The non-volatile solid state memory units may directly access the internal communication bus interface through a storage node communication bus, or request the storage node to access the bus interface. The non-volatile solid state memory unit contains an embedded CPU, solid state storage controller, and a quantity of solid state mass storage, e.g., between 2-32 terabytes ('TB') in some embodiments. An embedded volatile storage medium, such as DRAM, and an energy reserve apparatus are included in the non-volatile solid state memory unit. In some embodiments, the energy reserve apparatus is a capacitor, super-capacitor, or battery that enables transferring a subset of DRAM contents to a stable storage medium in the case of power loss. In some embodiments, the non-volatile solid state memory unit is constructed with a storage class memory, such as phase change or magnetoresistive random access memory ('MRAM') that substitutes for DRAM and enables a reduced power hold-up apparatus.

One of many features of the storage nodes and non-volatile solid state storage is the ability to proactively rebuild data in a storage cluster. The storage nodes and non-volatile solid state storage can determine when a storage node or non-volatile solid state storage in the storage cluster is unreachable, independent of whether there is an attempt to

13

read data involving that storage node or non-volatile solid state storage. The storage nodes and non-volatile solid state storage then cooperate to recover and rebuild the data in at least partially new locations. This constitutes a proactive rebuild, in that the system rebuilds data without waiting until the data is needed for a read access initiated from a client system employing the storage cluster. These and further details of the storage memory and operation thereof are discussed below.

FIG. 2A is a perspective view of a storage cluster 161, with multiple storage nodes 150 and internal solid-state memory coupled to each storage node to provide network attached storage or storage area network, in accordance with some embodiments. A network attached storage, storage area network, or a storage cluster, or other storage memory, could include one or more storage clusters 161, each having one or more storage nodes 150, in a flexible and reconfigurable arrangement of both the physical components and the amount of storage memory provided thereby. The storage cluster 161 is designed to fit in a rack, and one or more racks can be set up and populated as desired for the storage memory. The storage cluster 161 has a chassis 138 having multiple slots 142. It should be appreciated that chassis 138 may be referred to as a housing, enclosure, or rack unit. In one embodiment, the chassis 138 has fourteen slots 142, although other numbers of slots are readily devised. For example, some embodiments have four slots, eight slots, sixteen slots, thirty-two slots, or other suitable number of slots. Each slot 142 can accommodate one storage node 150 in some embodiments. Chassis 138 includes flaps 148 that can be utilized to mount the chassis 138 on a rack. Fans 144 provide air circulation for cooling of the storage nodes 150 and components thereof, although other cooling components could be used, or an embodiment could be devised without cooling components. A switch fabric 146 couples storage nodes 150 within chassis 138 together and to a network for communication to the memory. In an embodiment depicted in herein, the slots 142 to the left of the switch fabric 146 and fans 144 are shown occupied by storage nodes 150, while the slots 142 to the right of the switch fabric 146 and fans 144 are empty and available for insertion of storage node 150 for illustrative purposes. This configuration is one example, and one or more storage nodes 150 could occupy the slots 142 in various further arrangements. The storage node arrangements need not be sequential or adjacent in some embodiments. Storage nodes 150 are hot pluggable, meaning that a storage node 150 can be inserted into a slot 142 in the chassis 138, or removed from a slot 142, without stopping or powering down the system. Upon insertion or removal of storage node 150 from slot 142, the system automatically reconfigures in order to recognize and adapt to the change. Reconfiguration, in some embodiments, includes restoring redundancy and/or rebalancing data or load.

Each storage node 150 can have multiple components. In the embodiment shown here, the storage node 150 includes a printed circuit board 159 populated by a CPU 156, i.e., processor, a memory 154 coupled to the CPU 156, and a non-volatile solid state storage 152 coupled to the CPU 156, although other mountings and/or components could be used in further embodiments. The memory 154 has instructions which are executed by the CPU 156 and/or data operated on by the CPU 156. As further explained below, the non-volatile solid state storage 152 includes flash or, in further embodiments, other types of solid-state memory.

Referring to FIG. 2A, storage cluster 161 is scalable, meaning that storage capacity with non-uniform storage

14

sizes is readily added, as described above. One or more storage nodes 150 can be plugged into or removed from each chassis and the storage cluster self-configures in some embodiments. Plug-in storage nodes 150, whether installed in a chassis as delivered or later added, can have different sizes. For example, in one embodiment a storage node 150 can have any multiple of 4 TB, e.g., 8 TB, 12 TB, 16 TB, 32 TB, etc. In further embodiments, a storage node 150 could have any multiple of other storage amounts or capacities. Storage capacity of each storage node 150 is broadcast, and influences decisions of how to stripe the data. For maximum storage efficiency, an embodiment can self-configure as wide as possible in the stripe, subject to a predetermined requirement of continued operation with loss of up to one, or up to two, non-volatile solid state storage units 152 or storage nodes 150 within the chassis.

FIG. 2B is a block diagram showing a communications interconnect 173 and power distribution bus 172 coupling multiple storage nodes 150. Referring back to FIG. 2A, the communications interconnect 173 can be included in or implemented with the switch fabric 146 in some embodiments. Where multiple storage clusters 161 occupy a rack, the communications interconnect 173 can be included in or implemented with a top of rack switch, in some embodiments. As illustrated in FIG. 2B, storage cluster 161 is enclosed within a single chassis 138. External port 176 is coupled to storage nodes 150 through communications interconnect 173, while external port 174 is coupled directly to a storage node. External power port 178 is coupled to power distribution bus 172. Storage nodes 150 may include varying amounts and differing capacities of non-volatile solid state storage 152 as described with reference to FIG. 2A. In addition, one or more storage nodes 150 may be a compute only storage node as illustrated in FIG. 2B. Authorities 168 are implemented on the non-volatile solid state storages 152, for example as lists or other data structures stored in memory. In some embodiments the authorities are stored within the non-volatile solid state storage 152 and supported by software executing on a controller or other processor of the non-volatile solid state storage 152. In a further embodiment, authorities 168 are implemented on the storage nodes 150, for example as lists or other data structures stored in the memory 154 and supported by software executing on the CPU 156 of the storage node 150. Authorities 168 control how and where data is stored in the non-volatile solid state storages 152 in some embodiments. This control assists in determining which type of erasure coding scheme is applied to the data, and which storage nodes 150 have which portions of the data. Each authority 168 may be assigned to a non-volatile solid state storage 152. Each authority may control a range of inode numbers, segment numbers, or other data identifiers which are assigned to data by a file system, by the storage nodes 150, or by the non-volatile solid state storage 152, in various embodiments.

Every piece of data, and every piece of metadata, has redundancy in the system in some embodiments. In addition, every piece of data and every piece of metadata has an owner, which may be referred to as an authority. If that authority is unreachable, for example through failure of a storage node, there is a plan of succession for how to find that data or that metadata. In various embodiments, there are redundant copies of authorities 168. Authorities 168 have a relationship to storage nodes 150 and non-volatile solid state storage 152 in some embodiments. Each authority 168, covering a range of data segment numbers or other identifiers of the data, may be assigned to a specific non-volatile solid state storage 152. In some embodiments the authorities

168 for all of such ranges are distributed over the non-volatile solid state storages 152 of a storage cluster. Each storage node 150 has a network port that provides access to the non-volatile solid state storage(s) 152 of that storage node 150. Data can be stored in a segment, which is associated with a segment number and that segment number is an indirection for a configuration of a RAID (redundant array of independent disks) stripe in some embodiments. The assignment and use of the authorities 168 thus establishes an indirection to data. Indirection may be referred to as the ability to reference data indirectly, in this case via an authority 168, in accordance with some embodiments. A segment identifies a set of non-volatile solid state storage 152 and a local identifier into the set of non-volatile solid state storage 152 that may contain data. In some embodiments, the local identifier is an offset into the device and may be reused sequentially by multiple segments. In other embodiments the local identifier is unique for a specific segment and never reused. The offsets in the non-volatile solid state storage 152 are applied to locating data for writing to or reading from the non-volatile solid state storage 152 (in the form of a RAID stripe). Data is striped across multiple units of non-volatile solid state storage 152, which may include or be different from the non-volatile solid state storage 152 having the authority 168 for a particular data segment.

If there is a change in where a particular segment of data is located, e.g., during a data move or a data reconstruction, the authority 168 for that data segment should be consulted, at that non-volatile solid state storage 152 or storage node 150 having that authority 168. In order to locate a particular piece of data, embodiments calculate a hash value for a data segment or apply an inode number or a data segment number. The output of this operation points to a non-volatile solid state storage 152 having the authority 168 for that particular piece of data. In some embodiments there are two stages to this operation. The first stage maps an entity identifier (ID), e.g., a segment number, inode number, or directory number to an authority identifier. This mapping may include a calculation such as a hash or a bit mask. The second stage is mapping the authority identifier to a particular non-volatile solid state storage 152, which may be done through an explicit mapping. The operation is repeatable, so that when the calculation is performed, the result of the calculation repeatably and reliably points to a particular non-volatile solid state storage 152 having that authority 168. The operation may include the set of reachable storage nodes as input. If the set of reachable non-volatile solid state storage units changes the optimal set changes. In some embodiments, the persisted value is the current assignment (which is always true) and the calculated value is the target assignment the cluster will attempt to reconfigure towards. This calculation may be used to determine the optimal non-volatile solid state storage 152 for an authority in the presence of a set of non-volatile solid state storage 152 that are reachable and constitute the same cluster. The calculation also determines an ordered set of peer non-volatile solid state storage 152 that will also record the authority to non-volatile solid state storage mapping so that the authority may be determined even if the assigned non-volatile solid state storage is unreachable. A duplicate or substitute authority 168 may be consulted if a specific authority 168 is unavailable in some embodiments.

With reference to FIGS. 2A and 2B, two of the many tasks of the CPU 156 on a storage node 150 are to break up write data, and reassemble read data. When the system has determined that data is to be written, the authority 168 for that

data is located as above. When the segment ID for data is already determined the request to write is forwarded to the non-volatile solid state storage 152 currently determined to be the host of the authority 168 determined from the segment. The host CPU 156 of the storage node 150, on which the non-volatile solid state storage 152 and corresponding authority 168 reside, then breaks up or shards the data and transmits the data out to various non-volatile solid state storage 152. The transmitted data is written as a data stripe in accordance with an erasure coding scheme. In some embodiments, data is requested to be pulled, and in other embodiments, data is pushed. In reverse, when data is read, the authority 168 for the segment ID containing the data is located as described above. The host CPU 156 of the storage node 150 on which the non-volatile solid state storage 152 and corresponding authority 168 reside requests the data from the non-volatile solid state storage and corresponding storage nodes pointed to by the authority. In some embodiments the data is read from flash storage as a data stripe. The host CPU 156 of storage node 150 then reassembles the read data, correcting any errors (if present) according to the appropriate erasure coding scheme, and forwards the reassembled data to the network. In further embodiments, some or all of these tasks can be handled in the non-volatile solid state storage 152. In some embodiments, the segment host requests the data be sent to storage node 150 by requesting pages from storage and then sending the data to the storage node making the original request.

In some systems, for example in UNIX-style file systems, data is handled with an index node or inode, which specifies a data structure that represents an object in a file system. The object could be a file or a directory, for example. Metadata may accompany the object, as attributes such as permission data and a creation timestamp, among other attributes. A segment number could be assigned to all or a portion of such an object in a file system. In other systems, data segments are handled with a segment number assigned elsewhere. For purposes of discussion, the unit of distribution is an entity, and an entity can be a file, a directory, or a segment. That is, entities are units of data or metadata stored by a storage system. Entities are grouped into sets called authorities. Each authority has an authority owner, which is a storage node that has the exclusive right to update the entities in the authority. In other words, a storage node contains the authority, and that the authority, in turn, contains entities.

A segment is a logical container of data in accordance with some embodiments. A segment is an address space between medium address space and physical flash locations, i.e., the data segment number, are in this address space. Segments may also contain meta-data, which enable data redundancy to be restored (rewritten to different flash locations or devices) without the involvement of higher level software. In one embodiment, an internal format of a segment contains client data and medium mappings to determine the position of that data. Each data segment is protected, e.g., from memory and other failures, by breaking the segment into a number of data and parity shards, where applicable. The data and parity shards are distributed, i.e., striped, across non-volatile solid state storage 152 coupled to the host CPUs 156 (See FIGS. 2E and 2G) in accordance with an erasure coding scheme. Usage of the term segments refers to the container and its place in the address space of segments in some embodiments. Usage of the term stripe refers to the same set of shards as a segment and includes how the shards are distributed along with redundancy or parity information in accordance with some embodiments.

A series of address-space transformations takes place across an entire storage system. At the top are the directory entries (file names) which link to an inode. Inodes point into medium address space, where data is logically stored. Medium addresses may be mapped through a series of indirect mediums to spread the load of large files, or implement data services like deduplication or snapshots. Medium addresses may be mapped through a series of indirect mediums to spread the load of large files, or implement data services like deduplication or snapshots. Segment addresses are then translated into physical flash locations. Physical flash locations have an address range bounded by the amount of flash in the system in accordance with some embodiments. Medium addresses and segment addresses are logical containers, and in some embodiments use a 128 bit or larger identifier so as to be practically infinite, with a likelihood of reuse calculated as longer than the expected life of the system. Addresses from logical containers are allocated in a hierarchical fashion in some embodiments. Initially, each non-volatile solid state storage unit **152** may be assigned a range of address space. Within this assigned range, the non-volatile solid state storage **152** is able to allocate addresses without synchronization with other non-volatile solid state storage **152**.

Data and metadata is stored by a set of underlying storage layouts that are optimized for varying workload patterns and storage devices. These layouts incorporate multiple redundancy schemes, compression formats and index algorithms. Some of these layouts store information about authorities and authority masters, while others store file metadata and file data. The redundancy schemes include error correction codes that tolerate corrupted bits within a single storage device (such as a NAND flash chip), erasure codes that tolerate the failure of multiple storage nodes, and replication schemes that tolerate data center or regional failures. In some embodiments, low density parity check ('LDPC') code is used within a single storage unit. Reed-Solomon encoding is used within a storage cluster, and mirroring is used within a storage grid in some embodiments. Metadata may be stored using an ordered log structured index (such as a Log Structured Merge Tree), and large data may not be stored in a log structured layout.

In order to maintain consistency across multiple copies of an entity, the storage nodes agree implicitly on two things through calculations: (1) the authority that contains the entity, and (2) the storage node that contains the authority. The assignment of entities to authorities can be done by pseudo randomly assigning entities to authorities, by splitting entities into ranges based upon an externally produced key, or by placing a single entity into each authority. Examples of pseudorandom schemes are linear hashing and the Replication Under Scalable Hashing ('RUSH') family of hashes, including Controlled Replication Under Scalable Hashing ('CRUSH'). In some embodiments, pseudo-random assignment is utilized only for assigning authorities to nodes because the set of nodes can change. The set of authorities cannot change so any subjective function may be applied in these embodiments. Some placement schemes automatically place authorities on storage nodes, while other placement schemes rely on an explicit mapping of authorities to storage nodes. In some embodiments, a pseudorandom scheme is utilized to map from each authority to a set of candidate authority owners. A pseudorandom data distribution function related to CRUSH may assign authorities to storage nodes and create a list of where the authorities are assigned. Each storage node has a copy of the pseudorandom data distribution function, and can arrive at the same calcu-

lation for distributing, and later finding or locating an authority. Each of the pseudorandom schemes requires the reachable set of storage nodes as input in some embodiments in order to conclude the same target nodes. Once an entity has been placed in an authority, the entity may be stored on physical devices so that no expected failure will lead to unexpected data loss. In some embodiments, rebalancing algorithms attempt to store the copies of all entities within an authority in the same layout and on the same set of machines.

Examples of expected failures include device failures, stolen machines, datacenter fires, and regional disasters, such as nuclear or geological events. Different failures lead to different levels of acceptable data loss. In some embodiments, a stolen storage node impacts neither the security nor the reliability of the system, while depending on system configuration, a regional event could lead to no loss of data, a few seconds or minutes of lost updates, or even complete data loss.

In the embodiments, the placement of data for storage redundancy is independent of the placement of authorities for data consistency. In some embodiments, storage nodes that contain authorities do not contain any persistent storage. Instead, the storage nodes are connected to non-volatile solid state storage units that do not contain authorities. The communications interconnect between storage nodes and non-volatile solid state storage units consists of multiple communication technologies and has non-uniform performance and fault tolerance characteristics. In some embodiments, as mentioned above, non-volatile solid state storage units are connected to storage nodes via PCI express, storage nodes are connected together within a single chassis using Ethernet backplane, and chassis are connected together to form a storage cluster. Storage clusters are connected to clients using Ethernet or fiber channel in some embodiments. If multiple storage clusters are configured into a storage grid, the multiple storage clusters are connected using the Internet or other long-distance networking links, such as a "metro scale" link or private link that does not traverse the internet.

Authority owners have the exclusive right to modify entities, to migrate entities from one non-volatile solid state storage unit to another non-volatile solid state storage unit, and to add and remove copies of entities. This allows for maintaining the redundancy of the underlying data. When an authority owner fails, is going to be decommissioned, or is overloaded, the authority is transferred to a new storage node. Transient failures make it non-trivial to ensure that all non-faulty machines agree upon the new authority location. The ambiguity that arises due to transient failures can be achieved automatically by a consensus protocol such as Paxos, hot-warm failover schemes, via manual intervention by a remote system administrator, or by a local hardware administrator (such as by physically removing the failed machine from the cluster, or pressing a button on the failed machine). In some embodiments, a consensus protocol is used, and failover is automatic. If too many failures or replication events occur in too short a time period, the system goes into a self-preservation mode and halts replication and data movement activities until an administrator intervenes in accordance with some embodiments.

As authorities are transferred between storage nodes and authority owners update entities in their authorities, the system transfers messages between the storage nodes and non-volatile solid state storage units. With regard to persistent messages, messages that have different purposes are of different types. Depending on the type of the message, the

system maintains different ordering and durability guarantees. As the persistent messages are being processed, the messages are temporarily stored in multiple durable and non-durable storage hardware technologies. In some embodiments, messages are stored in RAM, NVRAM and on NAND flash devices, and a variety of protocols are used in order to make efficient use of each storage medium. Latency-sensitive client requests may be persisted in replicated NVRAM, and then later NAND, while background rebalancing operations are persisted directly to NAND.

Persistent messages are persistently stored prior to being transmitted. This allows the system to continue to serve client requests despite failures and component replacement. Although many hardware components contain unique identifiers that are visible to system administrators, manufacturer, hardware supply chain and ongoing monitoring quality control infrastructure, applications running on top of the infrastructure address virtualize addresses. These virtualized addresses do not change over the lifetime of the storage system, regardless of component failures and replacements. This allows each component of the storage system to be replaced over time without reconfiguration or disruptions of client request processing, i.e., the system supports non-disruptive upgrades.

In some embodiments, the virtualized addresses are stored with sufficient redundancy. A continuous monitoring system correlates hardware and software status and the hardware identifiers. This allows detection and prediction of failures due to faulty components and manufacturing details. The monitoring system also enables the proactive transfer of authorities and entities away from impacted devices before failure occurs by removing the component from the critical path in some embodiments.

FIG. 2C is a multiple level block diagram, showing contents of a storage node 150 and contents of a non-volatile solid state storage 152 of the storage node 150. Data is communicated to and from the storage node 150 by a network interface controller ('NIC') 202 in some embodiments. Each storage node 150 has a CPU 156, and one or more non-volatile solid state storage 152, as discussed above. Moving down one level in FIG. 2C, each non-volatile solid state storage 152 has a relatively fast non-volatile solid state memory, such as nonvolatile random access memory ('NVRAM') 204, and flash memory 206. In some embodiments, NVRAM 204 may be a component that does not require program/erase cycles (DRAM, MRAM, PCM), and can be a memory that can support being written vastly more often than the memory is read from. Moving down another level in FIG. 2C, the NVRAM 204 is implemented in one embodiment as high speed volatile memory, such as dynamic random access memory (DRAM) 216, backed up by energy reserve 218. Energy reserve 218 provides sufficient electrical power to keep the DRAM 216 powered long enough for contents to be transferred to the flash memory 206 in the event of power failure. In some embodiments, energy reserve 218 is a capacitor, super-capacitor, battery, or other device, that supplies a suitable supply of energy sufficient to enable the transfer of the contents of DRAM 216 to a stable storage medium in the case of power loss. The flash memory 206 is implemented as multiple flash dies 222, which may be referred to as packages of flash dies 222 or an array of flash dies 222. It should be appreciated that the flash dies 222 could be packaged in any number of ways, with a single die per package, multiple dies per package (i.e. multichip packages), in hybrid packages, as bare dies on a printed circuit board or other substrate, as encapsulated dies, etc. In the embodiment shown, the non-volatile solid state

storage 152 has a controller 212 or other processor, and an input output (I/O) port 210 coupled to the controller 212. I/O port 210 is coupled to the CPU 156 and/or the network interface controller 202 of the flash storage node 150. Flash input output (I/O) port 220 is coupled to the flash dies 222, and a direct memory access unit (DMA) 214 is coupled to the controller 212, the DRAM 216 and the flash dies 222. In the embodiment shown, the I/O port 210, controller 212, DMA unit 214 and flash I/O port 220 are implemented on a programmable logic device ('PLD') 208, e.g., a field programmable gate array (FPGA). In this embodiment, each flash die 222 has pages, organized as sixteen kB (kilobyte) pages 224, and a register 226 through which data can be written to or read from the flash die 222. In further embodiments, other types of solid-state memory are used in place of, or in addition to flash memory illustrated within flash die 222.

Storage clusters 161, in various embodiments as disclosed herein, can be contrasted with storage arrays in general. The storage nodes 150 are part of a collection that creates the storage cluster 161. Each storage node 150 owns a slice of data and computing required to provide the data. Multiple storage nodes 150 cooperate to store and retrieve the data. Storage memory or storage devices, as used in storage arrays in general, are less involved with processing and manipulating the data. Storage memory or storage devices in a storage array receive commands to read, write, or erase data. The storage memory or storage devices in a storage array are not aware of a larger system in which they are embedded, or what the data means. Storage memory or storage devices in storage arrays can include various types of storage memory, such as RAM, solid state drives, hard disk drives, etc. The storage units 152 described herein have multiple interfaces active simultaneously and serving multiple purposes. In some embodiments, some of the functionality of a storage node 150 is shifted into a storage unit 152, transforming the storage unit 152 into a combination of storage unit 152 and storage node 150. Placing computing (relative to storage data) into the storage unit 152 places this computing closer to the data itself. The various system embodiments have a hierarchy of storage node layers with different capabilities. By contrast, in a storage array, a controller owns and knows everything about all of the data that the controller manages in a shelf or storage devices. In a storage cluster 161, as described herein, multiple controllers in multiple storage units 152 and/or storage nodes 150 cooperate in various ways (e.g., for erasure coding, data sharding, metadata communication and redundancy, storage capacity expansion or contraction, data recovery, and so on).

FIG. 2D shows a storage server environment, which uses embodiments of the storage nodes 150 and storage units 152 of FIGS. 2A-C. In this version, each storage unit 152 has a processor such as controller 212 (see FIG. 2C), an FPGA (field programmable gate array), flash memory 206, and NVRAM 204 (which is super-capacitor backed DRAM 216, see FIGS. 2B and 2C) on a PCIe (peripheral component interconnect express) board in a chassis 138 (see FIG. 2A). The storage unit 152 may be implemented as a single board containing storage, and may be the largest tolerable failure domain inside the chassis. In some embodiments, up to two storage units 152 may fail and the device will continue with no data loss.

The physical storage is divided into named regions based on application usage in some embodiments. The NVRAM 204 is a contiguous block of reserved memory in the storage unit 152 DRAM 216, and is backed by NAND flash. NVRAM 204 is logically divided into multiple memory

regions written for two as spool (e.g., spool_region). Space within the NVRAM 204 spools is managed by each authority 168 independently. Each device provides an amount of storage space to each authority 168. That authority 168 further manages lifetimes and allocations within that space. 5 Examples of a spool include distributed transactions or notions. When the primary power to a storage unit 152 fails, onboard super-capacitors provide a short duration of power hold up. During this holdup interval, the contents of the NVRAM 204 are flushed to flash memory 206. On the next 10 power-on, the contents of the NVRAM 204 are recovered from the flash memory 206.

As for the storage unit controller, the responsibility of the logical “controller” is distributed across each of the blades containing authorities 168. This distribution of logical control is shown in FIG. 2D as a host controller 242, mid-tier controller 244 and storage unit controller(s) 246. Management of the control plane and the storage plane are treated independently, although parts may be physically co-located on the same blade. Each authority 168 effectively serves as 15 an independent controller. Each authority 168 provides its own data and metadata structures, its own background workers, and maintains its own lifecycle.

FIG. 2E is a blade 252 hardware block diagram, showing a control plane 254, compute and storage planes 256, 258, 25 and authorities 168 interacting with underlying physical resources, using embodiments of the storage nodes 150 and storage units 152 of FIGS. 2A-C in the storage server environment of FIG. 2D. The control plane 254 is partitioned into a number of authorities 168 which can use the compute resources in the compute plane 256 to run on any of the blades 252. The storage plane 258 is partitioned into a set of devices, each of which provides access to flash 206 and NVRAM 204 resources.

In the compute and storage planes 256, 258 of FIG. 2E, 35 the authorities 168 interact with the underlying physical resources (i.e., devices). From the point of view of an authority 168, its resources are striped over all of the physical devices. From the point of view of a device, it provides resources to all authorities 168, irrespective of 40 where the authorities happen to run. Each authority 168 has allocated or has been allocated one or more partitions 260 of storage memory in the storage units 152, e.g. partitions 260 in flash memory 206 and NVRAM 204. Each authority 168 uses those allocated partitions 260 that belong to it, for writing or reading user data. Authorities can be associated with differing amounts of physical storage of the system. For example, one authority 168 could have a larger number of partitions 260 or larger sized partitions 260 in one or more storage units 152 than one or more other authorities 168.

FIG. 2F depicts elasticity software layers in blades 252 of a storage cluster, in accordance with some embodiments. In the elasticity structure, elasticity software is symmetric, i.e., each blade’s compute module 270 runs the three identical layers of processes depicted in FIG. 2F. Storage managers 274 execute read and write requests from other blades 252 for data and metadata stored in local storage unit 152 NVRAM 204 and flash 206. Authorities 168 fulfill client requests by issuing the necessary reads and writes to the blades 252 on whose storage units 152 the corresponding data or metadata resides. Endpoints 272 parse client connection requests received from switch fabric 146 supervisory software, relay the client connection requests to the authorities 168 responsible for fulfillment, and relay the authorities’ 168 responses to clients. The symmetric three-layer structure enables the storage system’s high degree of concurrency. Elasticity scales out efficiently and reliably in

these embodiments. In addition, elasticity implements a unique scale-out technique that balances work evenly across all resources regardless of client access pattern, and maximizes concurrency by eliminating much of the need for inter-blade coordination that typically occurs with conventional distributed locking.

Still referring to FIG. 2F, authorities 168 running in the compute modules 270 of a blade 252 perform the internal operations required to fulfill client requests. One feature of elasticity is that authorities 168 are stateless, i.e., they cache active data and metadata in their own blades’ 252 DRAMs for fast access, but the authorities store every update in their NVRAM 204 partitions on three separate blades 252 until the update has been written to flash 206. All the storage 10 system writes to NVRAM 204 are in triplicate to partitions on three separate blades 252 in some embodiments. With triple-mirrored NVRAM 204 and persistent storage protected by parity and Reed-Solomon RAID checksums, the storage system can survive concurrent failure of two blades 252 with no loss of data, metadata, or access to either.

Because authorities 168 are stateless, they can migrate between blades 252. Each authority 168 has a unique identifier. NVRAM 204 and flash 206 partitions are associated with authorities’ 168 identifiers, not with the blades 252 on which they are running in some. Thus, when an authority 168 migrates, the authority 168 continues to manage the same storage partitions from its new location. When a new blade 252 is installed in an embodiment of the storage cluster, the system automatically rebalances load by: partitioning the new blade’s 252 storage for use by the system’s authorities 168, migrating selected authorities 168 to the new blade 252, starting endpoints 272 on the new blade 252 and including them in the switch fabric’s 146 client connection distribution algorithm.

From their new locations, migrated authorities 168 persist the contents of their NVRAM 204 partitions on flash 206, process read and write requests from other authorities 168, and fulfill the client requests that endpoints 272 direct to them. Similarly, if a blade 252 fails or is removed, the system redistributes its authorities 168 among the system’s remaining blades 252. The redistributed authorities 168 continue to perform their original functions from their new locations.

FIG. 2G depicts authorities 168 and storage resources in blades 252 of a storage cluster, in accordance with some embodiments. Each authority 168 is exclusively responsible for a partition of the flash 206 and NVRAM 204 on each blade 252. The authority 168 manages the content and integrity of its partitions independently of other authorities 168. Authorities 168 compress incoming data and preserve it temporarily in their NVRAM 204 partitions, and then consolidate, RAID-protect, and persist the data in segments of the storage in their flash 206 partitions. As the authorities 168 write data to flash 206, storage managers 274 perform the necessary flash translation to optimize write performance and maximize media longevity. In the background, authorities 168 “garbage collect,” or reclaim space occupied by data that clients have made obsolete by overwriting the data. It should be appreciated that since authorities’ 168 partitions are disjoint, there is no need for distributed locking to execute client and writes or to perform background functions.

The embodiments described herein may utilize various software, communication and/or networking protocols. In addition, the configuration of the hardware and/or software may be adjusted to accommodate various protocols. For example, the embodiments may utilize Active Directory,

which is a database based system that provides authentication, directory, policy, and other services in a WINDOWS' environment. In these embodiments, LDAP (Lightweight Directory Access Protocol) is one example application protocol for querying and modifying items in directory service providers such as Active Directory. In some embodiments, a network lock manager ('NLM') is utilized as a facility that works in cooperation with the Network File System ('NFS') to provide a System V style of advisory file and record locking over a network. The Server Message Block ('SMB') protocol, one version of which is also known as Common Internet File System ('CIFS'), may be integrated with the storage systems discussed herein. SMP operates as an application-layer network protocol typically used for providing shared access to files, printers, and serial ports and miscellaneous communications between nodes on a network. SMB also provides an authenticated inter-process communication mechanism. AMAZON™ S3 (Simple Storage Service) is a web service offered by Amazon Web Services, and the systems described herein may interface with Amazon S3 through web services interfaces (REST (representational state transfer), SOAP (simple object access protocol), and BitTorrent). A RESTful API (application programming interface) breaks down a transaction to create a series of small modules. Each module addresses a particular underlying part of the transaction. The control or permissions provided with these embodiments, especially for object data, may include utilization of an access control list ('ACL'). The ACL is a list of permissions attached to an object and the ACL specifies which users or system processes are granted access to objects, as well as what operations are allowed on given objects. The systems may utilize Internet Protocol version 6 ('IPv6'), as well as IPv4, for the communications protocol that provides an identification and location system for computers on networks and routes traffic across the Internet. The routing of packets between networked systems may include Equal-cost multi-path routing ('ECMP'), which is a routing strategy where next-hop packet forwarding to a single destination can occur over multiple "best paths" which tie for top place in routing metric calculations. Multi-path routing can be used in conjunction with most routing protocols because it is a per-hop decision limited to a single router. The software may support Multi-tenancy, which is an architecture in which a single instance of a software application serves multiple customers. Each customer may be referred to as a tenant. Tenants may be given the ability to customize some parts of the application, but may not customize the application's code, in some embodiments. The embodiments may maintain audit logs. An audit log is a document that records an event in a computing system. In addition to documenting what resources were accessed, audit log entries typically include destination and source addresses, a timestamp, and user login information for compliance with various regulations. The embodiments may support various key management policies, such as encryption key rotation. In addition, the system may support dynamic root passwords or some variation dynamically changing passwords.

FIG. 3A sets forth a diagram of a storage system 306 that is coupled for data communications with a cloud services provider 302 in accordance with some embodiments of the present disclosure. Although depicted in less detail, the storage system 306 depicted in FIG. 3A may be similar to the storage systems described above with reference to FIGS. 1A-1D and FIGS. 2A-2G. In some embodiments, the storage system 306 depicted in FIG. 3A may be embodied as a storage system that includes imbalanced active/active con-

trollers, as a storage system that includes balanced active/active controllers, as a storage system that includes active/active controllers where less than all of each controller's resources are utilized such that each controller has reserve resources that may be used to support failover, as a storage system that includes fully active/active controllers, as a storage system that includes dataset-segregated controllers, as a storage system that includes dual-layer architectures with front-end controllers and back-end integrated storage controllers, as a storage system that includes scale-out clusters of dual-controller arrays, as well as combinations of such embodiments.

In the example depicted in FIG. 3A, the storage system 306 is coupled to the cloud services provider 302 via a data communications link 304. The data communications link 304 may be embodied as a dedicated data communications link, as a data communications pathway that is provided through the use of one or data communications networks such as a wide area network ('WAN') or local area network ('LAN'), or as some other mechanism capable of transporting digital information between the storage system 306 and the cloud services provider 302. Such a data communications link 304 may be fully wired, fully wireless, or some aggregation of wired and wireless data communications pathways. In such an example, digital information may be exchanged between the storage system 306 and the cloud services provider 302 via the data communications link 304 using one or more data communications protocols. For example, digital information may be exchanged between the storage system 306 and the cloud services provider 302 via the data communications link 304 using the handheld device transfer protocol ('HDTP'), hypertext transfer protocol ('HTTP'), internet protocol ('IP'), real-time transfer protocol ('RTP'), transmission control protocol ('TCP'), user datagram protocol ('UDP'), wireless application protocol ('WAP'), or other protocol.

The cloud services provider 302 depicted in FIG. 3A may be embodied, for example, as a system and computing environment that provides services to users of the cloud services provider 302 through the sharing of computing resources via the data communications link 304. The cloud services provider 302 may provide on-demand access to a shared pool of configurable computing resources such as computer networks, servers, storage, applications, and services, and so on. The shared pool of configurable resources may be rapidly provisioned and released to a user of the cloud services provider 302 with minimal management effort. Generally, the user of the cloud services provider 302 is unaware of the exact computing resources utilized by the cloud services provider 302 to provide the services. Although in many cases such a cloud services provider 302 may be accessible via the Internet, readers of skill in the art will recognize that any system that abstracts the use of shared resources to provide services to a user through any data communications link may be considered a cloud services provider 302.

In the example depicted in FIG. 3A, the cloud services provider 302 may be configured to provide a variety of services to the storage system 306 and users of the storage system 306 through the implementation of various service models. For example, the cloud services provider 302 may be configured to provide services to the storage system 306 and users of the storage system 306 through the implementation of an infrastructure as a service ('IaaS') service model where the cloud services provider 302 offers computing infrastructure such as virtual machines and other resources as a service to subscribers. In addition, the cloud services

provider 302 may be configured to provide services to the storage system 306 and users of the storage system 306 through the implementation of a platform as a service ('PaaS') service model where the cloud services provider 302 offers a development environment to application developers. Such a development environment may include, for example, an operating system, programming-language execution environment, database, web server, or other components that may be utilized by application developers to develop and run software solutions on a cloud platform. Furthermore, the cloud services provider 302 may be configured to provide services to the storage system 306 and users of the storage system 306 through the implementation of a software as a service ('SaaS') service model where the cloud services provider 302 offers application software, databases, as well as the platforms that are used to run the applications to the storage system 306 and users of the storage system 306, providing the storage system 306 and users of the storage system 306 with on-demand software and eliminating the need to install and run the application on local computers, which may simplify maintenance and support of the application. The cloud services provider 302 may be further configured to provide services to the storage system 306 and users of the storage system 306 through the implementation of an authentication as a service ('AaaS') service model where the cloud services provider 302 offers authentication services that can be used to secure access to applications, data sources, or other resources. The cloud services provider 302 may also be configured to provide services to the storage system 306 and users of the storage system 306 through the implementation of a storage as a service model where the cloud services provider 302 offers access to its storage infrastructure for use by the storage system 306 and users of the storage system 306. Readers will appreciate that the cloud services provider 302 may be configured to provide additional services to the storage system 306 and users of the storage system 306 through the implementation of additional service models, as the service models described above are included only for explanatory purposes and in no way represent a limitation of the services that may be offered by the cloud services provider 302 or a limitation as to the service models that may be implemented by the cloud services provider 302.

In the example depicted in FIG. 3A, the cloud services provider 302 may be embodied, for example, as a private cloud, as a public cloud, or as a combination of a private cloud and public cloud. In an embodiment in which the cloud services provider 302 is embodied as a private cloud, the cloud services provider 302 may be dedicated to providing services to a single organization rather than providing services to multiple organizations. In an embodiment where the cloud services provider 302 is embodied as a public cloud, the cloud services provider 302 may provide services to multiple organizations. Public cloud and private cloud deployment models may differ and may come with various advantages and disadvantages. For example, because a public cloud deployment involves the sharing of a computing infrastructure across different organization, such a deployment may not be ideal for organizations with security concerns, mission-critical workloads, uptime requirements demands, and so on. While a private cloud deployment can address some of these issues, a private cloud deployment may require on-premises staff to manage the private cloud. In still alternative embodiments, the cloud services provider 302 may be embodied as a mix of a private and public cloud services with a hybrid cloud deployment.

Although not explicitly depicted in FIG. 3A, readers will appreciate that additional hardware components and additional software components may be necessary to facilitate the delivery of cloud services to the storage system 306 and users of the storage system 306. For example, the storage system 306 may be coupled to (or even include) a cloud storage gateway. Such a cloud storage gateway may be embodied, for example, as hardware-based or software-based appliance that is located on premise with the storage system 306. Such a cloud storage gateway may operate as a bridge between local applications that are executing on the storage system 306 and remote, cloud-based storage that is utilized by the storage system 306. Through the use of a cloud storage gateway, organizations may move primary iSCSI or NAS to the cloud services provider 302, thereby enabling the organization to save space on their on-premises storage systems. Such a cloud storage gateway may be configured to emulate a disk array, a block-based device, a file server, or other storage system that can translate the SCSI commands, file server commands, or other appropriate command into REST-space protocols that facilitate communications with the cloud services provider 302.

In order to enable the storage system 306 and users of the storage system 306 to make use of the services provided by the cloud services provider 302, a cloud migration process may take place during which data, applications, or other elements from an organization's local systems (or even from another cloud environment) are moved to the cloud services provider 302. In order to successfully migrate data, applications, or other elements to the cloud services provider's 302 environment, middleware such as a cloud migration tool may be utilized to bridge gaps between the cloud services provider's 302 environment and an organization's environment. Such cloud migration tools may also be configured to address potentially high network costs and long transfer times associated with migrating large volumes of data to the cloud services provider 302, as well as addressing security concerns associated with sensitive data to the cloud services provider 302 over data communications networks. In order to further enable the storage system 306 and users of the storage system 306 to make use of the services provided by the cloud services provider 302, a cloud orchestrator may also be used to arrange and coordinate automated tasks in pursuit of creating a consolidated process or workflow. Such a cloud orchestrator may perform tasks such as configuring various components, whether those components are cloud components or on-premises components, as well as managing the interconnections between such components. The cloud orchestrator can simplify the inter-component communication and connections to ensure that links are correctly configured and maintained.

In the example depicted in FIG. 3A, and as described briefly above, the cloud services provider 302 may be configured to provide services to the storage system 306 and users of the storage system 306 through the usage of a SaaS service model where the cloud services provider 302 offers application software, databases, as well as the platforms that are used to run the applications to the storage system 306 and users of the storage system 306, providing the storage system 306 and users of the storage system 306 with on-demand software and eliminating the need to install and run the application on local computers, which may simplify maintenance and support of the application. Such applications may take many forms in accordance with various embodiments of the present disclosure. For example, the cloud services provider 302 may be configured to provide access to data analytics applications to the storage system

306 and users of the storage system **306**. Such data analytics applications may be configured, for example, to receive telemetry data phoned home by the storage system **306**. Such telemetry data may describe various operating characteristics of the storage system **306** and may be analyzed, for example, to determine the health of the storage system **306**, to identify workloads that are executing on the storage system **306**, to predict when the storage system **306** will run out of various resources, to recommend configuration changes, hardware or software upgrades, workflow migrations, or other actions that may improve the operation of the storage system **306**.

The cloud services provider **302** may also be configured to provide access to virtualized computing environments to the storage system **306** and users of the storage system **306**. Such virtualized computing environments may be embodied, for example, as a virtual machine or other virtualized computer hardware platforms, virtual storage devices, virtualized computer network resources, and so on. Examples of such virtualized environments can include virtual machines that are created to emulate an actual computer, virtualized desktop environments that separate a logical desktop from a physical machine, virtualized file systems that allow uniform access to different types of concrete file systems, and many others.

For further explanation, FIG. 3B sets forth a diagram of a storage system **306** in accordance with some embodiments of the present disclosure. Although depicted in less detail, the storage system **306** depicted in FIG. 3B may be similar to the storage systems described above with reference to FIGS. 1A-1D and FIGS. 2A-2G as the storage system may include many of the components described above.

The storage system **306** depicted in FIG. 3B may include storage resources **308**, which may be embodied in many forms. For example, in some embodiments the storage resources **308** can include nano-RAM or another form of nonvolatile random access memory that utilizes carbon nanotubes deposited on a substrate. In some embodiments, the storage resources **308** may include 3D crosspoint non-volatile memory in which bit storage is based on a change of bulk resistance, in conjunction with a stackable cross-gridded data access array. In some embodiments, the storage resources **308** may include flash memory, including single-level cell ('SLC') NAND flash, multi-level cell ('MLC') NAND flash, triple-level cell ('TLC') NAND flash, quad-level cell ('QLC') NAND flash, and others. In some embodiments, the storage resources **308** may include non-volatile magnetoresistive random-access memory ('MRAM'), including spin transfer torque ('STT') MRAM, in which data is stored through the use of magnetic storage elements. In some embodiments, the example storage resources **308** may include non-volatile phase-change memory ('PCM') that may have the ability to hold multiple bits in a single cell as cells can achieve a number of distinct intermediary states. In some embodiments, the storage resources **308** may include quantum memory that allows for the storage and retrieval of photonic quantum information. In some embodiments, the example storage resources **308** may include resistive random-access memory ('ReRAM') in which data is stored by changing the resistance across a dielectric solid-state material. In some embodiments, the storage resources **308** may include storage class memory ('SCM') in which solid-state nonvolatile memory may be manufactured at a high density using some combination of sub-lithographic patterning techniques, multiple bits per cell, multiple layers of devices, and so on. Readers will appreciate that other forms of computer memories and storage devices

may be utilized by the storage systems described above, including DRAM, SRAM, EEPROM, universal memory, and many others. The storage resources **308** depicted in FIG. 3A may be embodied in a variety of form factors, including but not limited to, dual in-line memory modules ('DIMMs'), non-volatile dual in-line memory modules ('NVDIMMs'), M.2, U.2, and others.

The example storage system **306** depicted in FIG. 3B may implement a variety of storage architectures. For example, storage systems in accordance with some embodiments of the present disclosure may utilize block storage where data is stored in blocks, and each block essentially acts as an individual hard drive. Storage systems in accordance with some embodiments of the present disclosure may utilize object storage, where data is managed as objects. Each object may include the data itself, a variable amount of metadata, and a globally unique identifier, where object storage can be implemented at multiple levels (e.g., device level, system level, interface level). Storage systems in accordance with some embodiments of the present disclosure utilize file storage in which data is stored in a hierarchical structure. Such data may be saved in files and folders, and presented to both the system storing it and the system retrieving it in the same format.

The example storage system **306** depicted in FIG. 3B may be embodied as a storage system in which additional storage resources can be added through the use of a scale-up model, additional storage resources can be added through the use of a scale-out model, or through some combination thereof. In a scale-up model, additional storage may be added by adding additional storage devices. In a scale-out model, however, additional storage nodes may be added to a cluster of storage nodes, where such storage nodes can include additional processing resources, additional networking resources, and so on.

The storage system **306** depicted in FIG. 3B also includes communications resources **310** that may be useful in facilitating data communications between components within the storage system **306**, as well as data communications between the storage system **306** and computing devices that are outside of the storage system **306**. The communications resources **310** may be configured to utilize a variety of different protocols and data communication fabrics to facilitate data communications between components within the storage systems as well as computing devices that are outside of the storage system. For example, the communications resources **310** can include fibre channel ('FC') technologies such as FC fabrics and FC protocols that can transport SCSI commands over FC networks. The communications resources **310** can also include FC over ethernet ('FCoE') technologies through which FC frames are encapsulated and transmitted over Ethernet networks. The communications resources **310** can also include InfiniBand ('IB') technologies in which a switched fabric topology is utilized to facilitate transmissions between channel adapters. The communications resources **310** can also include NVM Express ('NVMe') technologies and NVMe over fabrics ('NVMeoF') technologies through which non-volatile storage media attached via a PCI express ('PCIe') bus may be accessed. The communications resources **310** can also include mechanisms for accessing storage resources **308** within the storage system **306** utilizing serial attached SCSI ('SAS'), serial ATA ('SATA') bus interfaces for connecting storage resources **308** within the storage system **306** to host bus adapters within the storage system **306**, internet small computer systems interface ('iSCSI') technologies to provide block-level access to storage resources **308** within the

storage system **306**, and other communications resources that that may be useful in facilitating data communications between components within the storage system **306**, as well as data communications between the storage system **306** and computing devices that are outside of the storage system **306**.

The storage system **306** depicted in FIG. 3B also includes processing resources **312** that may be useful in useful in executing computer program instructions and performing other computational tasks within the storage system **306**. The processing resources **312** may include one or more application-specific integrated circuits ('ASICs') that are customized for some particular purpose as well as one or more central processing units ('CPUs'). The processing resources **312** may also include one or more digital signal processors ('DSPs'), one or more field-programmable gate arrays ('FPGAs'), one or more systems on a chip ('SoCs'), or other form of processing resources **312**. The storage system **306** may utilize the storage resources **312** to perform a variety of tasks including, but not limited to, supporting the execution of software resources **314** that will be described in greater detail below.

The storage system **306** depicted in FIG. 3B also includes software resources **314** that, when executed by processing resources **312** within the storage system **306**, may perform various tasks. The software resources **314** may include, for example, one or more modules of computer program instructions that when executed by processing resources **312** within the storage system **306** are useful in carrying out various data protection techniques to preserve the integrity of data that is stored within the storage systems. Readers will appreciate that such data protection techniques may be carried out, for example, by system software executing on computer hardware within the storage system, by a cloud services provider, or in other ways. Such data protection techniques can include, for example, data archiving techniques that cause data that is no longer actively used to be moved to a separate storage device or separate storage system for long-term retention, data backup techniques through which data stored in the storage system may be copied and stored in a distinct location to avoid data loss in the event of equipment failure or some other form of catastrophe with the storage system, data replication techniques through which data stored in the storage system is replicated to another storage system such that the data may be accessible via multiple storage systems, data snapshotting techniques through which the state of data within the storage system is captured at various points in time, data and database cloning techniques through which duplicate copies of data and databases may be created, and other data protection techniques. Through the use of such data protection techniques, business continuity and disaster recovery objectives may be met as a failure of the storage system may not result in the loss of data stored in the storage system.

The software resources **314** may also include software that is useful in implementing software-defined storage ('SDS'). In such an example, the software resources **314** may include one or more modules of computer program instructions that, when executed, are useful in policy-based provisioning and management of data storage that is independent of the underlying hardware. Such software resources **314** may be useful in implementing storage virtualization to separate the storage hardware from the software that manages the storage hardware.

The software resources **314** may also include software that is useful in facilitating and optimizing I/O operations that are directed to the storage resources **308** in the storage

system **306**. For example, the software resources **314** may include software modules that perform carry out various data reduction techniques such as, for example, data compression, data deduplication, and others. The software resources **314** may include software modules that intelligently group together I/O operations to facilitate better usage of the underlying storage resource **308**, software modules that perform data migration operations to migrate from within a storage system, as well as software modules that perform other functions. Such software resources **314** may be embodied as one or more software containers or in many other ways.

Readers will appreciate that the presence of such software resources **314** may provide for an improved user experience of the storage system **306**, an expansion of functionality supported by the storage system **306**, and many other benefits. Consider the specific example of the software resources **314** carrying out data backup techniques through which data stored in the storage system may be copied and stored in a distinct location to avoid data loss in the event of equipment failure or some other form of catastrophe. In such an example, the systems described herein may more reliably (and with less burden placed on the user) perform backup operations relative to interactive backup management systems that require high degrees of user interactivity, offer less robust automation and feature sets, and so on.

The storage systems described above may carry out intelligent data backup techniques through which data stored in the storage system may be copied and stored in a distinct location to avoid data loss in the event of equipment failure or some other form of catastrophe. For example, the storage systems described above may be configured to examine each backup to avoid restoring the storage system to an undesirable state. Consider an example in which malware infects the storage system. In such an example, the storage system may include software resources **314** that can scan each backup to identify backups that were captured before the malware infected the storage system and those backups that were captured after the malware infected the storage system. In such an example, the storage system may restore itself from a backup that does not include the malware—or at least not restore the portions of a backup that contained the malware. In such an example, the storage system may include software resources **314** that can scan each backup to identify the presences of malware (or a virus, or some other undesirable), for example, by identifying write operations that were serviced by the storage system and originated from a network subnet that is suspected to have delivered the malware, by identifying write operations that were serviced by the storage system and originated from a user that is suspected to have delivered the malware, by identifying write operations that were serviced by the storage system and examining the content of the write operation against fingerprints of the malware, and in many other ways.

Readers will further appreciate that the backups (often in the form of one or more snapshots) may also be utilized to perform rapid recovery of the storage system. Consider an example in which the storage system is infected with ransomware that locks users out of the storage system. In such an example, software resources **314** within the storage system may be configured to detect the presence of ransomware and may be further configured to restore the storage system to a point-in-time, using the retained backups, prior to the point-in-time at which the ransomware infected the storage system. In such an example, the presence of ransomware may be explicitly detected through the use of software tools utilized by the system, through the use of a

key (e.g., a USB drive) that is inserted into the storage system, or in a similar way. Likewise, the presence of ransomware may be inferred in response to system activity meeting a predetermined fingerprint such as, for example, no reads or writes coming into the system for a predetermined period of time.

Readers will appreciate that the various components depicted in FIG. 3B may be grouped into one or more optimized computing packages as converged infrastructures. Such converged infrastructures may include pools of computers, storage and networking resources that can be shared by multiple applications and managed in a collective manner using policy-driven processes. Such converged infrastructures may minimize compatibility issues between various components within the storage system 306 while also reducing various costs associated with the establishment and operation of the storage system 306. Such converged infrastructures may be implemented with a converged infrastructure reference architecture, with standalone appliances, with a software driven hyper-converged approach (e.g., hyper-converged infrastructures), or in other ways.

Readers will appreciate that the storage system 306 depicted in FIG. 3B may be useful for supporting various types of software applications. For example, the storage system 306 may be useful in supporting artificial intelligence ('AI') applications, database applications, DevOps projects, electronic design automation tools, event-driven software applications, high performance computing applications, simulation applications, high-speed data capture and analysis applications, machine learning applications, media production applications, media serving applications, picture archiving and communication systems ('PACS') applications, software development applications, virtual reality applications, augmented reality applications, and many other types of applications by providing storage resources to such applications.

The storage systems described above may operate to support a wide variety of applications. In view of the fact that the storage systems include compute resources, storage resources, and a wide variety of other resources, the storage systems may be well suited to support applications that are resource intensive such as, for example, AI applications. Such AI applications may enable devices to perceive their environment and take actions that maximize their chance of success at some goal. Examples of such AI applications can include IBM Watson, Microsoft Oxford, Google DeepMind, Baidu Minwa, and others. The storage systems described above may also be well suited to support other types of applications that are resource intensive such as, for example, machine learning applications. Machine learning applications may perform various types of data analysis to automate analytical model building. Using algorithms that iteratively learn from data, machine learning applications can enable computers to learn without being explicitly programmed.

In addition to the resources already described, the storage systems described above may also include graphics processing units ('GPUs'), occasionally referred to as visual processing unit ('VPUs'). Such GPUs may be embodied as specialized electronic circuits that rapidly manipulate and alter memory to accelerate the creation of images in a frame buffer intended for output to a display device. Such GPUs may be included within any of the computing devices that are part of the storage systems described above, including as one of many individually scalable components of a storage system, where other examples of individually scalable components of such storage system can include storage components, memory components, compute components (e.g.,

CPUs, FPGAs, ASICs), networking components, software components, and others. In addition to GPUs, the storage systems described above may also include neural network processors ('NNPs') for use in various aspects of neural network processing. Such NNPs may be used in place of (or in addition to) GPUs and may be also be independently scalable.

As described above, the storage systems described herein may be configured to support artificial intelligence applications, machine learning applications, big data analytics applications, and many other types of applications. The rapid growth in these sort of applications is being driven by three technologies: deep learning (DL), GPU processors, and Big Data. Deep learning is a computing model that makes use of massively parallel neural networks inspired by the human brain. Instead of experts handcrafting software, a deep learning model writes its own software by learning from lots of examples. A GPU is a modern processor with thousands of cores, well-suited to run algorithms that loosely represent the parallel nature of the human brain.

Advances in deep neural networks have ignited a new wave of algorithms and tools for data scientists to tap into their data with artificial intelligence (AI). With improved algorithms, larger data sets, and various frameworks (including open-source software libraries for machine learning across a range of tasks), data scientists are tackling new use cases like autonomous driving vehicles, natural language processing and understanding, computer vision, machine reasoning, strong AI, and many others. Applications of such techniques may include: machine and vehicular object detection, identification and avoidance; visual recognition, classification and tagging; algorithmic financial trading strategy performance management; simultaneous localization and mapping; predictive maintenance of high-value machinery; prevention against cyber security threats, expertise automation; image recognition and classification; question answering; robotics; text analytics (extraction, classification) and text generation and translation; and many others. Applications of AI techniques has materialized in a wide array of products include, for example, Amazon Echo's speech recognition technology that allows users to talk to their machines, Google Translate™ which allows for machine-based language translation, Spotify's Discover Weekly that provides recommendations on new songs and artists that a user may like based on the user's usage and traffic analysis, Quill's text generation offering that takes structured data and turns it into narrative stories, Chatbots that provide real-time, contextually specific answers to questions in a dialog format, and many others. Furthermore, AI may impact a wide variety of industries and sectors. For example, AI solutions may be used in healthcare to take clinical notes, patient files, research data, and other inputs to generate potential treatment options for doctors to explore. Likewise, AI solutions may be used by retailers to personalize consumer recommendations based on a person's digital footprint of behaviors, profile data, or other data.

Training deep neural networks, however, requires both high quality input data and large amounts of computation. GPUs are massively parallel processors capable of operating on large amounts of data simultaneously. When combined into a multi-GPU cluster, a high throughput pipeline may be required to feed input data from storage to the compute engines. Deep learning is more than just constructing and training models. There also exists an entire data pipeline that must be designed for the scale, iteration, and experimentation necessary for a data science team to succeed.

Data is the heart of modern AI and deep learning algorithms. Before training can begin, one problem that must be addressed revolves around collecting the labeled data that is crucial for training an accurate AI model. A full scale AI deployment may be required to continuously collect, clean, transform, label, and store large amounts of data. Adding additional high quality data points directly translates to more accurate models and better insights. Data samples may undergo a series of processing steps including, but not limited to: 1) ingesting the data from an external source into the training system and storing the data in raw form, 2) cleaning and transforming the data in a format convenient for training, including linking data samples to the appropriate label, 3) exploring parameters and models, quickly testing with a smaller dataset, and iterating to converge on the most promising models to push into the production cluster, 4) executing training phases to select random batches of input data, including both new and older samples, and feeding those into production GPU servers for computation to update model parameters, and 5) evaluating including using a holdback portion of the data not used in training in order to evaluate model accuracy on the holdout data. This lifecycle may apply for any type of parallelized machine learning, not just neural networks, or deep learning. For example, standard machine learning frameworks may rely on CPUs instead of GPUs but the data ingest and training workflows may be the same. Readers will appreciate that a single shared storage data hub creates a coordination point throughout the lifecycle without the need for extra data copies among the ingest, preprocessing, and training stages. Rarely is the ingested data used for only one purpose, and shared storage gives the flexibility to train multiple different models or apply traditional analytics to the data.

Readers will appreciate that each stage in the AI data pipeline may have varying requirements from the data hub (e.g., the storage system or collection of storage systems). Scale-out storage systems must deliver uncompromising performance for all manner of access types and patterns—from small, metadata-heavy to large files, from random to sequential access patterns, and from low to high concurrency. The storage systems described above may serve as an ideal AI data hub as the systems may service unstructured workloads. In the first stage, data is ideally ingested and stored on to the same data hub that following stages will use, in order to avoid excess data copying. The next two steps can be done on a standard compute server that optionally includes a GPU, and then in the fourth and last stage, full training production jobs are run on powerful GPU-accelerated servers. Often, there is a production pipeline alongside an experimental pipeline operating on the same dataset. Further, the GPU-accelerated servers can be used independently for different models or joined together to train on one larger model, even spanning multiple systems for distributed training. If the shared storage tier is slow, then data must be copied to local storage for each phase, resulting in wasted time staging data onto different servers. The ideal data hub for the AI training pipeline delivers performance similar to data stored locally on the server node while also having the simplicity and performance to enable all pipeline stages to operate concurrently.

A data scientist works to improve the usefulness of the trained model through a wide variety of approaches: more data, better data, smarter training, and deeper models. In many cases, there will be teams of data scientists sharing the same datasets and working in parallel to produce new and improved training models. Often, there is a team of data scientists working within these phases concurrently on the

same shared datasets. Multiple, concurrent workloads of data processing, experimentation, and full-scale training layer the demands of multiple access patterns on the storage tier. In other words, storage cannot just satisfy large file reads, but must contend with a mix of large and small file reads and writes. Finally, with multiple data scientists exploring datasets and models, it may be critical to store data in its native format to provide flexibility for each user to transform, clean, and use the data in a unique way. The storage systems described above may provide a natural shared storage home for the dataset, with data protection redundancy (e.g., by using RAID6) and the performance necessary to be a common access point for multiple developers and multiple experiments. Using the storage systems described above may avoid the need to carefully copy subsets of the data for local work, saving both engineering and GPU-accelerated servers use time. These copies become a constant and growing tax as the raw data set and desired transformations constantly update and change.

Readers will appreciate that a fundamental reason why deep learning has seen a surge in success is the continued improvement of models with larger data set sizes. In contrast, classical machine learning algorithms, like logistic regression, stop improving in accuracy at smaller data set sizes. As such, the separation of compute resources and storage resources may also allow independent scaling of each tier, avoiding many of the complexities inherent in managing both together. As the data set size grows or new data sets are considered, a scale out storage system must be able to expand easily. Similarly, if more concurrent training is required, additional GPUs or other compute resources can be added without concern for their internal storage. Furthermore, the storage systems described above may make building, operating, and growing an AI system easier due to the random read bandwidth provided by the storage systems, the ability to of the storage systems to randomly read small files (50 KB) high rates (meaning that no extra effort is required to aggregate individual data points to make larger, storage-friendly files), the ability of the storage systems to scale capacity and performance as either the dataset grows or the throughput requirements grow, the ability of the storage systems to support files or objects, the ability of the storage systems to tune performance for large or small files (i.e., no need for the user to provision filesystems), the ability of the storage systems to support non-disruptive upgrades of hardware and software even during production model training, and for many other reasons.

Small file performance of the storage tier may be critical as many types of inputs, including text, audio, or images will be natively stored as small files. If the storage tier does not handle small files well, an extra step will be required to pre-process and group samples into larger files. Storage, built on top of spinning disks, that relies on SSD as a caching tier, may fall short of the performance needed. Because training with random input batches results in more accurate models, the entire data set must be accessible with full performance. SSD caches only provide high performance for a small subset of the data and will be ineffective at hiding the latency of spinning drives.

Although the preceding paragraphs discuss deep learning applications, readers will appreciate that the storage systems described herein may also be part of a distributed deep learning ('DDL') platform to support the execution of DDL algorithms. Distributed deep learning may can be used to significantly accelerate deep learning with distributed computing on GPUs (or other form of accelerator or computer program instruction executor), such that parallelism can be

achieved. In addition, the output of training machine learning and deep learning models, such as a fully trained machine learning model, may be used for a variety of purposes and in conjunction with other tools. For example, trained machine learning models may be used in conjunction with tools like Core ML to integrate a broad variety of machine learning model types into an application. In fact, trained models may be run through Core ML converter tools and inserted into a custom application that can be deployed on compatible devices. The storage systems described above may also be paired with other technologies such as TensorFlow, an open-source software library for dataflow programming across a range of tasks that may be used for machine learning applications such as neural networks, to facilitate the development of such machine learning models, applications, and so on.

Readers will further appreciate that the systems described above may be deployed in a variety of ways to support the democratization of AI, as AI becomes more available for mass consumption. The democratization of AI may include, for example, the ability to offer AI as a Platform-as-a-Service, the growth of Artificial general intelligence offerings, the proliferation of Autonomous level 4 and Autonomous level 5 vehicles, the availability of autonomous mobile robots, the development of conversational AI platforms, and many others. For example, the systems described above may be deployed in cloud environments, edge environments, or other environments that are useful in supporting the democratization of AI. As part of the democratization of AI, a movement may occur from narrow AI that consists of highly scoped machine learning solutions that target a particular task to artificial general intelligence where the use of machine learning is expanded to handle a broad range of use cases that could essentially perform any intelligent task that a human could perform and could learn dynamically, much like a human.

The storage systems described above may also be used in a neuromorphic computing environment. Neuromorphic computing is a form of computing that mimics brain cells. To support neuromorphic computing, an architecture of interconnected "neurons" replace traditional computing models with low-powered signals that go directly between neurons for more efficient computation. Neuromorphic computing may make use of very-large-scale integration (VLSI) systems containing electronic analog circuits to mimic neuro-biological architectures present in the nervous system, as well as analog, digital, mixed-mode analog/digital VLSI, and software systems that implement models of neural systems for perception, motor control, or multisensory integration.

Readers will appreciate that the storage systems described above may be configured to support the storage or use of (among other types of data) blockchains. Such blockchains may be embodied as a continuously growing list of records, called blocks, which are linked and secured using cryptography. Each block in a blockchain may contain a hash pointer as a link to a previous block, a timestamp, transaction data, and so on. Blockchains may be designed to be resistant to modification of the data and can serve as an open, distributed ledger that can record transactions between two parties efficiently and in a verifiable and permanent way. This makes blockchains potentially suitable for the recording of events, medical records, and other records management activities, such as identity management, transaction processing, and others. In addition to supporting the storage and use of blockchain technologies, the storage systems described above may also support the storage and use of

derivative items such as, for example, open source blockchains and related tools that are part of the IBM™ Hyperledger project, permissioned blockchains in which a certain number of trusted parties are allowed to access the block chain, blockchain products that enable developers to build their own distributed ledger projects, and others. Readers will appreciate that blockchain technologies may impact a wide variety of industries and sectors. For example, blockchain technologies may be used in real estate transactions as blockchain based contracts whose use can eliminate the need for 3rd parties and enable self-executing actions when conditions are met. Likewise, universal health records can be created by aggregating and placing a person's health history onto a blockchain ledger for any healthcare provider, or permissioned health care providers, to access and update.

Readers will appreciate that the usage of blockchains is not limited to financial transactions, contracts, and the like. In fact, blockchains may be leveraged to enable the decentralized aggregation, ordering, timestamping, and archiving of any type of information, including structured data, correspondence, documentation, or other data. Through the usage of blockchains, participants can provably and permanently agree on exactly what data was entered, when and by whom, without relying on a trusted intermediary. For example, SAP's recently launched blockchain platform, which supports MultiChain and Hyperledger Fabric, targets a broad range of supply chain and other non-financial applications.

One way to use a blockchain for recording data is to embed each piece of data directly inside a transaction. Every blockchain transaction may be digitally signed by one or more parties, replicated to a plurality of nodes, ordered, and timestamped by the chain's consensus algorithm, and stored permanently in a tamper-proof way. Any data within the transaction will therefore be stored identically but independently by every node, along with a proof of who wrote it and when. The chain's users are able to retrieve this information at any future time. This type of storage may be referred to as on-chain storage. On-chain storage may not be particularly practical, however, when attempting to store a very large dataset. As such, in accordance with embodiments of the present disclosure, blockchains and the storage systems described herein may be leveraged to support on-chain storage of data as well as off-chain storage of data.

Off-chain storage of data can be implemented in a variety of ways and can occur when the data itself is not stored within the blockchain. For example, in one embodiment, a hash function may be utilized and the data itself may be fed into the hash function to generate a hash value. In such an example, the hashes of large pieces of data may be embedded within transactions, instead of the data itself. Each hash may serve as a commitment to its input data, with the data itself being stored outside of the blockchain. Readers will appreciate that any blockchain participant that needs an off-chain piece of data cannot reproduce the data from its hash, but if the data can be retrieved in some other way, then the on-chain hash serves to confirm who created it and when. Just like regular on-chain data, the hash may be embedded inside a digitally signed transaction, which was included in the chain by consensus.

Readers will appreciate that, in other embodiments, alternatives to blockchains may be used to facilitate the decentralized storage of information. For example, one alternative to a blockchain that may be used is a blockweave. While conventional blockchains store every transaction to achieve validation, a blockweave permits secure decentralization without the usage of the entire chain, thereby enabling low

cost on-chain storage of data. Such blockweaves may utilize a consensus mechanism that is based on proof of access (PoA) and proof of work (PoW). While typical PoW systems only depend on the previous block in order to generate each successive block, the PoA algorithm may incorporate data from a randomly chosen previous block. Combined with the blockweave data structure, miners do not need to store all blocks (forming a blockchain), but rather can store any previous blocks forming a weave of blocks (a blockweave). This enables increased levels of scalability, speed and low-cost and reduces the cost of data storage in part because miners need not store all blocks, thereby resulting in a substantial reduction in the amount of electricity that is consumed during the mining process because, as the network expands, electricity consumption decreases because a blockweave demands less and less hashing power for consensus as data is added to the system. Furthermore, blockweaves may be deployed on a decentralized storage network in which incentives are created to encourage rapid data sharing. Such decentralized storage networks may also make use of blockshadowing techniques, where nodes only send a minimal block “shadow” to other nodes that allows peers to reconstruct a full block, instead of transmitting the full block itself.

The storage systems described above may, either alone or in combination with other computing devices, be used to support in-memory computing applications. In memory computing involves the storage of information in RAM that is distributed across a cluster of computers. In-memory computing helps business customers, including retailers, banks, and utilities, to quickly detect patterns, analyze massive data volumes on the fly, and perform their operations quickly. Readers will appreciate that the storage systems described above, especially those that are configurable with customizable amounts of processing resources, storage resources, and memory resources (e.g., those systems in which blades that contain configurable amounts of each type of resource), may be configured in a way so as to provide an infrastructure that can support in-memory computing. Likewise, the storage systems described above may include component parts (e.g., NVDIMMs, 3D crosspoint storage that provide fast random access memory that is persistent) that can actually provide for an improved in-memory computing environment as compared to in-memory computing environments that rely on RAM distributed across dedicated servers.

In some embodiments, the storage systems described above may be configured to operate as a hybrid in-memory computing environment that includes a universal interface to all storage media (e.g., RAM, flash storage, 3D crosspoint storage). In such embodiments, users may have no knowledge regarding the details of where their data is stored but they can still use the same full, unified API to address data. In such embodiments, the storage system may (in the background) move data to the fastest layer available—including intelligently placing the data in dependence upon various characteristics of the data or in dependence upon some other heuristic. In such an example, the storage systems may even make use of existing products such as Apache Ignite and GridGain to move data between the various storage layers, or the storage systems may make use of custom software to move data between the various storage layers. The storage systems described herein may implement various optimizations to improve the performance of in-memory computing such as, for example, having computations occur as close to the data as possible.

Readers will further appreciate that in some embodiments, the storage systems described above may be paired with other resources to support the applications described above. For example, one infrastructure could include primary compute in the form of servers and workstations which specialize in using General-purpose computing on graphics processing units (‘GPGPU’) to accelerate deep learning applications that are interconnected into a computation engine to train parameters for deep neural networks. Each system may have Ethernet external connectivity, InfiniBand external connectivity, some other form of external connectivity, or some combination thereof. In such an example, the GPUs can be grouped for a single large training or used independently to train multiple models. The infrastructure could also include a storage system such as those described above to provide, for example, a scale-out all-flash file or object store through which data can be accessed via high-performance protocols such as NFS, S3, and so on. The infrastructure can also include, for example, redundant top-of-rack Ethernet switches connected to storage and compute via ports in MLAG port channels for redundancy. The infrastructure could also include additional compute in the form of whitebox servers, optionally with GPUs, for data ingestion, pre-processing, and model debugging. Readers will appreciate that additional infrastructures are also be possible.

Readers will appreciate that the systems described above may be better suited for the applications described above relative to other systems that may include, for example, a distributed direct-attached storage (DDAS) solution deployed in server nodes. Such DDAS solutions may be built for handling large, less sequential accesses but may be less able to handle small, random accesses. Readers will further appreciate that the storage systems described above may be utilized to provide a platform for the applications described above that is preferable to the utilization of cloud-based resources as the storage systems may be included in an on-site or in-house infrastructure that is more secure, more locally and internally managed, more robust in feature sets and performance, or otherwise preferable to the utilization of cloud-based resources as part of a platform to support the applications described above. For example, services built on platforms such as IBM’s Watson may require a business enterprise to distribute individual user information, such as financial transaction information or identifiable patient records, to other institutions. As such, cloud-based offerings of AI as a service may be less desirable than internally managed and offered AI as a service that is supported by storage systems such as the storage systems described above, for a wide array of technical reasons as well as for various business reasons.

Readers will appreciate that the storage systems described above, either alone or in coordination with other computing machinery may be configured to support other AI related tools. For example, the storage systems may make use of tools like ONXX or other open neural network exchange formats that make it easier to transfer models written in different AI frameworks. Likewise, the storage systems may be configured to support tools like Amazon’s Gluon that allow developers to prototype, build, and train deep learning models. In fact, the storage systems described above may be part of a larger platform, such as IBM™ Cloud Private for Data, that includes integrated data science, data engineering and application building services. Such platforms may seamlessly collect, organize, secure, and analyze data across an enterprise, as well as simplify hybrid data management,

unified data governance and integration, data science and business analytics with a single solution.

Readers will further appreciate that the storage systems described above may also be deployed as an edge solution. Such an edge solution may be in place to optimize cloud computing systems by performing data processing at the edge of the network, near the source of the data. Edge computing can push applications, data, and computing power (i.e., services) away from centralized points to the logical extremes of a network. Through the use of edge solutions such as the storage systems described above, computational tasks may be performed using the compute resources provided by such storage systems, data may be storage using the storage resources of the storage system, and cloud-based services may be accessed through the use of various resources of the storage system (including networking resources). By performing computational tasks on the edge solution, storing data on the edge solution, and generally making use of the edge solution, the consumption of expensive cloud-based resources may be avoided and, in fact, performance improvements may be experienced relative to a heavier reliance on cloud-based resources.

While many tasks may benefit from the utilization of an edge solution, some particular uses may be especially suited for deployment in such an environment. For example, devices like drones, autonomous cars, robots, and others may require extremely rapid processing—so fast, in fact, that sending data up to a cloud environment and back to receive data processing support may simply be too slow. Likewise, machines like locomotives and gas turbines that generate large amounts of information through the use of a wide array of data-generating sensors may benefit from the rapid data processing capabilities of an edge solution. As an additional example, some IoT devices such as connected video cameras may not be well-suited for the utilization of cloud-based resources as it may be impractical (not only from a privacy perspective, security perspective, or a financial perspective) to send the data to the cloud simply because of the pure volume of data that is involved. As such, many tasks that really on data processing, storage, or communications may be better suited by platforms that include edge solutions such as the storage systems described above.

Consider a specific example of inventory management in a warehouse, distribution center, or similar location. A large inventory, warehousing, shipping, order-fulfillment, manufacturing, or other operation has a large amount of inventory on inventory shelves, and high resolution digital cameras that produce a firehose of large data. All of this data may be taken into an image processing system, which may reduce the amount of data to a firehose of small data. All of the small data may be stored on-premises in storage. The on-premises storage, at the edge of the facility, may be coupled to the cloud, for external reports, real-time control, and cloud storage. Inventory management may be performed with the results of the image processing, so that inventory can be tracked on the shelves and restocked, moved, shipped, modified with new products, or discontinued/obsolescent products deleted, etc. The above scenario is a prime candidate for an embodiment of the configurable processing and storage systems described above. A combination of compute-only blades and offload blades suited for the image processing, perhaps with deep learning on offload-FPGA or offload-custom blade(s) could take in the firehose of large data from all of the digital cameras, and produce the firehose of small data. All of the small data could then be stored by storage nodes, operating with storage units in whichever combination of types of storage

blades best handles the data flow. This is an example of storage and function acceleration and integration. Depending on external communication needs with the cloud, and external processing in the cloud, and depending on reliability of network connections and cloud resources, the system could be sized for storage and compute management with bursty workloads and variable conductivity reliability. Also, depending on other inventory management aspects, the system could be configured for scheduling and resource management in a hybrid edge/cloud environment.

The storage systems described above may alone, or in combination with other computing resources, serves as a network edge platform that combines compute resources, storage resources, networking resources, cloud technologies and network virtualization technologies, and so on. As part of the network, the edge may take on characteristics similar to other network facilities, from the customer premise and backhaul aggregation facilities to Points of Presence (PoPs) and regional data centers. Readers will appreciate that network workloads, such as Virtual Network Functions (VNFs) and others, will reside on the network edge platform. Enabled by a combination of containers and virtual machines, the network edge platform may rely on controllers and schedulers that are no longer geographically co-located with the data processing resources. The functions, as microservices, may split into control planes, user, and data planes, or even state machines, allowing for independent optimization and scaling techniques to be applied. Such user and data planes may be enabled through increased accelerators, both those residing in server platforms, such as FPGAs and Smart NICs, and through SDN-enabled merchant silicon and programmable ASICs.

The storage systems described above may also be optimized for use in big data analytics. Big data analytics may be generally described as the process of examining large and varied data sets to uncover hidden patterns, unknown correlations, market trends, customer preferences and other useful information that can help organizations make more-informed business decisions. Big data analytics applications enable data scientists, predictive modelers, statisticians, and other analytics professionals to analyze growing volumes of structured transaction data, plus other forms of data that are often left untapped by conventional business intelligence (BI) and analytics programs. As part of that process, semi-structured and unstructured data such as, for example, internet clickstream data, web server logs, social media content, text from customer emails and survey responses, mobile-phone call-detail records, IoT sensor data, and other data may be converted to a structured form. Big data analytics is a form of advanced analytics, which involves complex applications with elements such as predictive models, statistical algorithms and what-if analyses powered by high-performance analytics systems.

The storage systems described above may also support (including implementing as a system interface) applications that perform tasks in response to human speech. For example, the storage systems may support the execution intelligent personal assistant applications such as, for example, Amazon's Alexa, Apple Siri, Google Voice, Samsung Bixby, Microsoft Cortana, and others. While the examples described in the previous sentence make use of voice as input, the storage systems described above may also support chatbots, talkbots, chatterbots, or artificial conversational entities or other applications that are configured to conduct a conversation via auditory or textual methods. Likewise, the storage system may actually execute such an application to enable a user such as a system administrator

to interact with the storage system via speech. Such applications are generally capable of voice interaction, music playback, making to-do lists, setting alarms, streaming podcasts, playing audiobooks, and providing weather, traffic, and other real time information, such as news, although in 5 embodiments in accordance with the present disclosure, such applications may be utilized as interfaces to various system management operations.

The storage systems described above may also implement AI platforms for delivering on the vision of self-driving 10 storage. Such AI platforms may be configured to deliver global predictive intelligence by collecting and analyzing large amounts of storage system telemetry data points to enable effortless management, analytics, and support. In fact, such storage systems may be capable of predicting both 15 capacity and performance, as well as generating intelligent advice on workload deployment, interaction, and optimization. Such AI platforms may be configured to scan all incoming storage system telemetry data against a library of issue fingerprints to predict and resolve incidents in real-time, before they impact customer environments, and captures hundreds of variables related to performance that are used to forecast performance load.

The storage systems described above may support the serialized or simultaneous execution artificial intelligence 25 applications, machine learning applications, data analytics applications, data transformations, and other tasks that collectively may form an AI ladder. Such an AI ladder may effectively be formed by combining such elements to form a complete data science pipeline, where exist dependencies between elements of the AI ladder. For example, AI may require that some form of machine learning has taken place, machine learning may require that some form of analytics 30 has taken place, analytics may require that some form of data and information architecting has taken place, and so on. As such, each element may be viewed as a rung in an AI ladder that collectively can form a complete and sophisticated AI solution.

The storage systems described above may also, either alone or in combination with other computing environments, be used to deliver an AI everywhere experience where AI permeates wide and expansive aspects of business and life. For example, AI may play an important role in the delivery of deep learning solutions, deep reinforcement learning solutions, artificial general intelligence solutions, autonomous vehicles, cognitive computing solutions, commercial UAVs or drones, conversational user interfaces, enterprise taxonomies, ontology management solutions, machine learning solutions, smart dust, smart robots, smart workplaces, and many others. The storage systems described 40 above may also, either alone or in combination with other computing environments, be used to deliver a wide range of transparently immersive experiences where technology can introduce transparency between people, businesses, and things. Such transparently immersive experiences may be delivered as augmented reality technologies, connected homes, virtual reality technologies, brain—computer interfaces, human augmentation technologies, nanotube electronics, volumetric displays, 4D printing technologies, or others. The storage systems described above may also, either alone 50 or in combination with other computing environments, be used to support a wide variety of digital platforms. Such digital platforms can include, for example, 5G wireless systems and platforms, digital twin platforms, edge computing platforms, IoT platforms, quantum computing platforms, serverless PaaS, software-defined security, neuro-morphic computing platforms, and so on.

Readers will appreciate that some transparently immersive experiences may involve the use of digital twins of various “things” such as people, places, processes, systems, and so on. Such digital twins and other immersive technologies can alter the way that humans interact with technology, as conversational platforms, augmented reality, virtual reality, and mixed reality provide a more natural and immersive interaction with the digital world. In fact, digital twins may be linked with the real-world, perhaps even in real-time, to understand the state of a thing or system, respond to changes, and so on. Because digital twins consolidate massive amounts of information on individual assets and groups of assets (even possibly providing control of those assets), digital twins may communicate with each other to digital 15 factory models of multiple linked digital twins.

The storage systems described above may also be part of a multi-cloud environment in which multiple cloud computing and storage services are deployed in a single heterogeneous architecture. In order to facilitate the operation of such a multi-cloud environment, DevOps tools may be deployed to enable orchestration across clouds. Likewise, continuous development and continuous integration tools may be deployed to standardize processes around continuous integration and delivery, new feature rollout and provisioning cloud workloads. By standardizing these processes, a multi-cloud strategy may be implemented that enables the utilization of the best provider for each workload. Furthermore, application monitoring and visibility tools may be deployed to move application workloads around different clouds, identify performance issues, and perform other tasks. In addition, security and compliance tools may be deployed for to ensure compliance with security requirements, government regulations, and so on. Such a multi-cloud environment may also include tools for application delivery and smart workload management to ensure efficient application delivery and help direct workloads across the distributed and heterogeneous infrastructure, as well as tools that ease the deployment and maintenance of packaged and custom applications in the cloud and enable portability amongst clouds. The multi-cloud environment may similarly include tools for data portability.

The storage systems described above may be used as a part of a platform to enable the use of crypto-anchors that may be used to authenticate a product’s origins and contents to ensure that it matches a blockchain record associated with the product. Such crypto-anchors may take many forms including, for example, as edible ink, as a mobile sensor, as a microchip, and others. Similarly, as part of a suite of tools to secure data stored on the storage system, the storage systems described above may implement various encryption technologies and schemes, including lattice cryptography. Lattice cryptography can involve constructions of cryptographic primitives that involve lattices, either in the construction itself or in the security proof. Unlike public-key schemes such as the RSA, Diffie-Hellman, or Elliptic-Curve cryptosystems, which are easily attacked by a quantum computer, some lattice-based constructions appear to be resistant to attack by both classical and quantum computers.

A quantum computer is a device that performs quantum computing. Quantum computing is computing using quantum-mechanical phenomena, such as superposition and entanglement. Quantum computers differ from traditional computers that are based on transistors, as such traditional computers require that data be encoded into binary digits (bits), each of which is always in one of two definite states (0 or 1). In contrast to traditional computers, quantum computers use quantum bits, which can be in superpositions 65

of states. A quantum computer maintains a sequence of qubits, where a single qubit can represent a one, a zero, or any quantum superposition of those two qubit states. A pair of qubits can be in any quantum superposition of 4 states, and three qubits in any superposition of 8 states. A quantum computer with n qubits can generally be in an arbitrary superposition of up to 2^n different states simultaneously, whereas a traditional computer can only be in one of these states at any one time. A quantum Turing machine is a theoretical model of such a computer.

The storage systems described above may also be paired with FPGA-accelerated servers as part of a larger AI or ML infrastructure. Such FPGA-accelerated servers may reside near (e.g., in the same data center) the storage systems described above or even incorporated into an appliance that includes one or more storage systems, one or more FPGA-accelerated servers, networking infrastructure that supports communications between the one or more storage systems and the one or more FPGA-accelerated servers, as well as other hardware and software components. Alternatively, FPGA-accelerated servers may reside within a cloud computing environment that may be used to perform compute-related tasks for AI and ML jobs. Any of the embodiments described above may be used to collectively serve as a FPGA-based AI or ML platform. Readers will appreciate that, in some embodiments of the FPGA-based AI or ML platform, the FPGAs that are contained within the FPGA-accelerated servers may be reconfigured for different types of ML models (e.g., LSTMs, CNNs, GRUs). The ability to reconfigure the FPGAs that are contained within the FPGA-accelerated servers may enable the acceleration of a ML or AI application based on the most optimal numerical precision and memory model being used. Readers will appreciate that by treating the collection of FPGA-accelerated servers as a pool of FPGAs, any CPU in the data center may utilize the pool of FPGAs as a shared hardware microservice, rather than limiting a server to dedicated accelerators plugged into it.

The FPGA-accelerated servers and the GPU-accelerated servers described above may implement a model of computing where, rather than keeping a small amount of data in a CPU and running a long stream of instructions over it as occurred in more traditional computing models, the machine learning model and parameters are pinned into the high-bandwidth on-chip memory with lots of data streaming through the high-bandwidth on-chip memory. FPGAs may even be more efficient than GPUs for this computing model, as the FPGAs can be programmed with only the instructions needed to run this kind of computing model.

The storage systems described above may be configured to provide parallel storage, for example, through the use of a parallel file system such as BeeGFS. Such parallel file systems may include a distributed metadata architecture. For example, the parallel file system may include a plurality of metadata servers across which metadata is distributed, as well as components that include services for clients and storage servers. Through the use of a parallel file system, file contents may be distributed over a plurality of storage servers using striping and metadata may be distributed over a plurality of metadata servers on a directory level, with each server storing a part of the complete file system tree. Readers will appreciate that in some embodiments, the storage servers and metadata servers may run in userspace on top of an existing local file system. Furthermore, dedicated hardware is not required for client services, the metadata servers, or

the hardware servers as metadata servers, storage servers, and even the client services may be run on the same machines.

Readers will appreciate that, in part due to the emergence of many of the technologies discussed above including mobile devices, cloud services, social networks, big data analytics, and so on, an information technology platform may be needed to integrate all of these technologies and drive new business opportunities by quickly delivering revenue-generating products, services, and experiences—rather than merely providing the technology to automate internal business processes. Information technology organizations may need to balance resources and investments needed to keep core legacy systems up and running while also integrating technologies to build an information technology platform that can provide the speed and flexibility in areas such as, for example, exploiting big data, managing unstructured data, and working with cloud applications and services. One possible embodiment of such an information technology platform is a composable infrastructure that includes fluid resource pools, such as many of the systems described above that, can meet the changing needs of applications by allowing for the composition and recomposition of blocks of disaggregated compute, storage, and fabric infrastructure. Such a composable infrastructure can also include a single management interface to eliminate complexity and a unified API to discover, search, inventory, configure, provision, update, and diagnose the composable infrastructure.

For further explanation, FIG. 4 sets forth a block diagram illustrating a plurality of storage systems (402, 404, 406) that support a pod according to some embodiments of the present disclosure. Although depicted in less detail, the storage systems (402, 404, 406) depicted in FIG. 4 may be similar to the storage systems described above with reference to FIGS. 1A-1D, FIGS. 2A-2G, FIGS. 3A-3B, or any combination thereof. In fact, the storage systems (402, 404, 406) depicted in FIG. 4 may include the same, fewer, or additional components as the storage systems described above.

In the example depicted in FIG. 4, each of the storage systems (402, 404, 406) is depicted as having at least one computer processor (408, 410, 412), computer memory (414, 416, 418), and computer storage (420, 422, 424). Although in some embodiments the computer memory (414, 416, 418) and the computer storage (420, 422, 424) may be part of the same hardware devices, in other embodiments the computer memory (414, 416, 418) and the computer storage (420, 422, 424) may be part of different hardware devices. The distinction between the computer memory (414, 416, 418) and the computer storage (420, 422, 424) in this particular example may be that the computer memory (414, 416, 418) is physically proximate to the computer processors (408, 410, 412) and may store computer program instructions that are executed by the computer processors (408, 410, 412), while the computer storage (420, 422, 424) is embodied as non-volatile storage for storing user data, metadata describing the user data, and so on. Referring to the example above in FIG. 1A, for example, the computer processors (408, 410, 412) and computer memory (414, 416, 418) for a particular storage system (402, 404, 406) may reside within one of more of the controllers (110A-110D) while the attached storage devices (171A-171F) may serve as the computer storage (420, 422, 424) within a particular storage system (402, 404, 406).

In the example depicted in FIG. 4, the depicted storage systems (402, 404, 406) may attach to one or more pods

(430, 432) according to some embodiments of the present disclosure. Each of the pods (430, 432) depicted in FIG. 4 can include a dataset (426, 428). For example, a first pod (430) that three storage systems (402, 404, 406) have attached to includes a first dataset (426) while a second pod (432) that two storage systems (404, 406) have attached to includes a second dataset (428). In such an example, when a particular storage system attaches to a pod, the pod's dataset is copied to the particular storage system and then kept up to date as the dataset is modified. Storage systems can be removed from a pod, resulting in the dataset being no longer kept up to date on the removed storage system. In the example depicted in FIG. 4, any storage system which is active for a pod (it is an up-to-date, operating, non-faulted member of a non-faulted pod) can receive and process requests to modify or read the pod's dataset.

In the example depicted in FIG. 4, each pod (430, 432) may also include a set of managed objects and management operations, as well as a set of access operations to modify or read the dataset (426, 428) that is associated with the particular pod (430, 432). In such an example, the management operations may modify or query managed objects equivalently through any of the storage systems. Likewise, access operations to read or modify the dataset may operate equivalently through any of the storage systems. In such an example, while each storage system stores a separate copy of the dataset as a proper subset of the datasets stored and advertised for use by the storage system, the operations to modify managed objects or the dataset performed and completed through any one storage system are reflected in subsequent management objects to query the pod or subsequent access operations to read the dataset.

Readers will appreciate that pods may implement more capabilities than just a clustered synchronously replicated dataset. For example, pods can be used to implement tenants, whereby datasets are in some way securely isolated from each other. Pods can also be used to implement virtual arrays or virtual storage systems where each pod is presented as a unique storage entity on a network (e.g., a Storage Area Network, or Internet Protocol network) with separate addresses. In the case of a multi-storage-system pod implementing a virtual storage system, all physical storage systems associated with the pod may present themselves as in some way the same storage system (e.g., as if the multiple physical storage systems were no different than multiple network ports into a single storage system).

Readers will appreciate that pods may also be units of administration, representing a collection of volumes, file systems, object/analytic stores, snapshots, and other administrative entities, where making administrative changes (e.g., name changes, property changes, managing exports or permissions for some part of the pod's dataset), on any one storage system is automatically reflected to all active storage systems associated with the pod. In addition, pods could also be units of data collection and data analysis, where performance and capacity metrics are presented in ways that aggregate across all active storage systems for the pod, or that call out data collection and analysis separately for each pod, or perhaps presenting each attached storage system's contribution to the incoming content and performance for each a pod.

One model for pod membership may be defined as a list of storage systems, and a subset of that list where storage systems are considered to be in-sync for the pod. A storage system may be considered to be in-sync for a pod if it is at least within a recovery of having identical idle content for the last written copy of the dataset associated with the pod.

Idle content is the content after any in-progress modifications have completed with no processing of new modifications. Sometimes this is referred to as "crash recoverable" consistency. Recovery of a pod carries out the process of reconciling differences in applying concurrent updates to in-sync storage systems in the pod. Recovery can resolve any inconsistencies between storage systems in the completion of concurrent modifications that had been requested to various members of the pod but that were not signaled to any requestor as having completed successfully. Storage systems that are listed as pod members but that are not listed as in-sync for the pod can be described as "detached" from the pod. Storage systems that are listed as pod members, are in-sync for the pod, and are currently available for actively serving data for the pod are "online" for the pod.

Each storage system member of a pod may have its own copy of the membership, including which storage systems it last knew were in-sync, and which storage systems it last knew comprised the entire set of pod members. To be online for a pod, a storage system must consider itself to be in-sync for the pod and must be communicating with all other storage systems it considers to be in-sync for the pod. If a storage system can't be certain that it is in-sync and communicating with all other storage systems that are in-sync, then it must stop processing new incoming requests for the pod (or must complete them with an error or exception) until it can be certain that it is in-sync and communicating with all other storage systems that are in-sync. A first storage system may conclude that a second paired storage system should be detached, which will allow the first storage system to continue since it is now in-sync with all storage systems now in the list. But, the second storage system must be prevented from concluding, alternatively, that the first storage system should be detached and with the second storage system continuing operation. This would result in a "split brain" condition that can lead to irreconcilable datasets, dataset corruption, or application corruption, among other dangers.

The situation of needing to determine how to proceed when not communicating with paired storage systems can arise while a storage system is running normally and then notices lost communications, while it is currently recovering from some previous fault, while it is rebooting or resuming from a temporary power loss or recovered communication outage, while it is switching operations from one set of storage system controller to another set for whatever reason, or during or after any combination of these or other kinds of events. In fact, any time a storage system that is associated with a pod can't communicate with all known non-detached members, the storage system can either wait briefly until communications can be established, go offline and continue waiting, or it can determine through some means that it is safe to detach the non-communicating storage system without risk of incurring a split brain due to the non-communicating storage system concluding the alternative view, and then continue. If a safe detach can happen quickly enough, the storage system can remain online for the pod with little more than a short delay and with no resulting application outages for applications that can issue requests to the remaining online storage systems.

One example of this situation is when a storage system may know that it is out-of-date. That can happen, for example, when a first storage system is first added to a pod that is already associated with one or more storage systems, or when a first storage system reconnects to another storage system and finds that the other storage system had already marked the first storage system as detached. In this case, this

first storage system will simply wait until it connects to some other set of storage systems that are in-sync for the pod.

This model demands some degree of consideration for how storage systems are added to or removed from pods or from the in-sync pod members list. Since each storage system will have its own copy of the list, and since two independent storage systems can't update their local copy at exactly the same time, and since the local copy is all that is available on a reboot or in various fault scenarios, care must be taken to ensure that transient inconsistencies don't cause problems. For example, if one storage system is in-sync for a pod and a second storage system is added, then if the second storage system is updated to list both storage systems as in-sync first, then if there is a fault and a restart of both storage systems, the second might startup and wait to connect to the first storage system while the first might be unaware that it should or could wait for the second storage system. If the second storage system then responds to an inability to connect with the first storage system by going through a process to detach it, then it might succeed in completing a process that the first storage system is unaware of, resulting in a split brain. As such, it may be necessary to ensure that storage systems won't disagree inappropriately on whether they might opt to go through a detach process if they aren't communicating.

One way to ensure that storage systems won't disagree inappropriately on whether they might opt to go through a detach process if they aren't communicating is to ensure that when adding a new storage system to the in-sync member list for a pod, the new storage system first stores that it is a detached member (and perhaps that it is being added as an in-sync member). Then, the existing in-sync storage systems can locally store that the new storage system is an in-sync pod member before the new storage system locally stores that same fact. If there is a set of reboots or network outages prior to the new storage system storing its in-sync status, then the original storage systems may detach the new storage system due to non-communication, but the new storage system will wait. A reverse version of this change might be needed for removing a communicating storage system from a pod: first the storage system being removed stores that it is no longer in-sync, then the storage systems that will remain store that the storage system being removed is no longer in-sync, then all storage systems delete the storage system being removed from their pod membership lists. Depending on the implementation, an intermediate persisted detached state may not be necessary. Whether or not care is required in local copies of membership lists may depend on the model storage systems use for monitoring each other or for validating their membership. If a consensus model is used for both, or if an external system (or an external distributed or clustered system) is used to store and validate pod membership, then inconsistencies in locally stored membership lists may not matter.

When communications fail or one or several storage systems in a pod fail, or when a storage system starts up (or fails over to a secondary controller) and can't communicate with paired storage systems for a pod, and it is time for one or more storage systems to decide to detach one or more paired storage systems, some algorithm or mechanism must be employed to decide that it is safe to do so and to follow through on the detach. One means of resolving detaches is use a majority (or quorum) model for membership. With three storage systems, as long as two are communicating, they can agree to detach a third storage system that isn't communicating, but that third storage system cannot by itself choose to detach either of the other two. Confusion can

arise when storage system communication is inconsistent. For example, storage system A might be communicating with storage system B but not C, while storage system B might be communicating with both A and C. So, A and B could detach C, or B and C could detach A, but more communication between pod members may be needed to figure this out.

Care needs to be taken in a quorum membership model when adding and removing storage systems. For example, if a fourth storage system is added, then a "majority" of storage systems is at that point three. The transition from three storage systems (with two required for majority) to a pod including a fourth storage system (with three required for majority) may require something similar to the model described previously for carefully adding a storage system to the in-sync list. For example, the fourth storage system might start in an attaching state but not yet attached where it would never instigate a vote over quorum. Once in that state, the original three pod members could each be updated to be aware of the fourth member and the new requirement for a three storage system majority to detach a fourth. Removing a storage system from a pod might similarly move that storage system to a locally stored "detaching" state before updating other pod members. A variant scheme for this is to use a distributed consensus mechanism such as PAXOS or RAFT to implement any membership changes or to process detach requests.

Another means of managing membership transitions is to use an external system that is outside of the storage systems themselves to handle pod membership. In order to become online for a pod, a storage system must first contact the external pod membership system to verify that it is in-sync for the pod. Any storage system that is online for a pod should then remain in communication with the pod membership system and should wait or go offline if it loses communication. An external pod membership manager could be implemented as a highly available cluster using various cluster tools, such as Oracle RAC, Linux HA, VERITAS Cluster Server, IBM's HACMP, or others. An external pod membership manager could also use distributed configuration tools such as EtcD or Zookeeper, or a reliable distributed database such as Amazon's DynamoDB.

In the example depicted in FIG. 4, the depicted storage systems (402, 404, 406) may receive a request to read a portion of the dataset (426, 428) and process the request to read the portion of the dataset locally according to some embodiments of the present disclosure. Readers will appreciate that although requests to modify (e.g., a write operation) the dataset (426, 428) require coordination between the storage systems (402, 404, 406) in a pod, as the dataset (426, 428) should be consistent across all storage systems (402, 404, 406) in a pod, responding to a request to read a portion of the dataset (426, 428) does not require similar coordination between the storage systems (402, 404, 406). As such, a particular storage system that receives a read request may service the read request locally by reading a portion of the dataset (426, 428) that is stored within the storage system's storage devices, with no synchronous communication with other storage systems in the pod. Read requests received by one storage system for a replicated dataset in a replicated cluster are expected to avoid any communication in the vast majority of cases, at least when received by a storage system that is running within a cluster that is also running nominally. Such reads should normally be processed simply by reading from the local copy of a clustered dataset with no further interaction required with other storage systems in the cluster

Readers will appreciate that the storage systems may take steps to ensure read consistency such that a read request will return the same result regardless of which storage system processes the read request. For example, the resulting clustered dataset content for any set of updates received by any set of storage systems in the cluster should be consistent across the cluster, at least at any time updates are idle (all previous modifying operations have been indicated as complete and no new update requests have been received and processed in any way). More specifically, the instances of a clustered dataset across a set of storage systems can differ only as a result of updates that have not yet completed. This means, for example, that any two write requests which overlap in their volume block range, or any combination of a write request and an overlapping snapshot, compare-and-write, or virtual block range copy, must yield a consistent result on all copies of the dataset. Two operations should not yield a result as if they happened in one order on one storage system and a different order on another storage system in the replicated cluster.

Furthermore, read requests can be made time order consistent. For example, if one read request is received on a replicated cluster and completed and that read is then followed by another read request to an overlapping address range which is received by the replicated cluster and where one or both reads in any way overlap in time and volume address range with a modification request received by the replicated cluster (whether any of the reads or the modification are received by the same storage system or a different storage system in the replicated cluster), then if the first read reflects the result of the update then the second read should also reflect the results of that update, rather than possibly returning data that preceded the update. If the first read does not reflect the update, then the second read can either reflect the update or not. This ensures that between two read requests “time” for a data segment cannot roll backward.

In the example depicted in FIG. 4, the depicted storage systems (402, 404, 406) may also detect a disruption in data communications with one or more of the other storage systems and determine whether to the particular storage system should remain in the pod. A disruption in data communications with one or more of the other storage systems may occur for a variety of reasons. For example, a disruption in data communications with one or more of the other storage systems may occur because one of the storage systems has failed, because a network interconnect has failed, or for some other reason. An important aspect of synchronous replicated clustering is ensuring that any fault handling doesn’t result in unrecoverable inconsistencies, or any inconsistency in responses. For example, if a network fails between two storage systems, at most one of the storage systems can continue processing newly incoming I/O requests for a pod. And, if one storage system continues processing, the other storage system can’t process any new requests to completion, including read requests.

In the example depicted in FIG. 4, the depicted storage systems (402, 404, 406) may also determine whether the particular storage system should remain in the pod in response to detecting a disruption in data communications with one or more of the other storage systems. As mentioned above, to be ‘online’ as part of a pod, a storage system must consider itself to be in-sync for the pod and must be communicating with all other storage systems it considers to be in-sync for the pod. If a storage system can’t be certain that it is in-sync and communicating with all other storage systems that are in-sync, then it may stop processing new incoming requests to access the dataset (426, 428). As such,

the storage system may determine whether to the particular storage system should remain online as part of the pod, for example, by determining whether it can communicate with all other storage systems it considers to be in-sync for the pod (e.g., via one or more test messages), by determining whether the all other storage systems it considers to be in-sync for the pod also consider the storage system to be attached to the pod, through a combination of both steps where the particular storage system must confirm that it can communicate with all other storage systems it considers to be in-sync for the pod and that all other storage systems it considers to be in-sync for the pod also consider the storage system to be attached to the pod, or through some other mechanism.

In the example depicted in FIG. 4, the depicted storage systems (402, 404, 406) may also keep the dataset on the particular storage system accessible for management and dataset operations in response to determining that the particular storage system should remain in the pod. The storage system may keep the dataset (426, 428) on the particular storage system accessible for management and dataset operations, for example, by accepting requests to access the version of the dataset (426, 428) that is stored on the storage system and processing such requests, by accepting and processing management operations associated with the dataset (426, 428) that are issued by a host or authorized administrator, by accepting and processing management operations associated with the dataset (426, 428) that are issued by one of the other storage systems, or in some other way.

In the example depicted in FIG. 4, the depicted storage systems (402, 404, 406) may, however, make the dataset on the particular storage system inaccessible for management and dataset operations in response to determining that the particular storage system should not remain in the pod. The storage system may make the dataset (426, 428) on the particular storage system inaccessible for management and dataset operations, for example, by rejecting requests to access the version of the dataset (426, 428) that is stored on the storage system, by rejecting management operations associated with the dataset (426, 428) that are issued by a host or other authorized administrator, by rejecting management operations associated with the dataset (426, 428) that are issued by one of the other storage systems in the pod, or in some other way.

In the example depicted in FIG. 4, the depicted storage systems (402, 404, 406) may also detect that the disruption in data communications with one or more of the other storage systems has been repaired and make the dataset on the particular storage system accessible for management and dataset operations. The storage system may detect that the disruption in data communications with one or more of the other storage systems has been repaired, for example, by receiving a message from the one or more of the other storage systems. In response to detecting that the disruption in data communications with one or more of the other storage systems has been repaired, the storage system may make the dataset (426, 428) on the particular storage system accessible for management and dataset operations once the previously detached storage system has been resynchronized with the storage systems that remained attached to the pod.

In the example depicted in FIG. 4, the depicted storage systems (402, 404, 406) may also go offline from the pod such that the particular storage system no longer allows management and dataset operations. The depicted storage systems (402, 404, 406) may go offline from the pod such that the particular storage system no longer allows manage-

ment and dataset operations for a variety of reasons. For example, the depicted storage systems (402, 404, 406) may also go offline from the pod due to some fault with the storage system itself, because an update or some other maintenance is occurring on the storage system, due to communications faults, or for many other reasons. In such an example, the depicted storage systems (402, 404, 406) may subsequently update the dataset on the particular storage system to include all updates to the dataset since the particular storage system went offline and go back online with the pod such that the particular storage system allows management and dataset operations, as will be described in greater detail in the resynchronization sections included below.

In the example depicted in FIG. 4, the depicted storage systems (402, 404, 406) may also identify a target storage system for asynchronously receiving the dataset, where the target storage system is not one of the plurality of storage systems across which the dataset is synchronously replicated. Such a target storage system may represent, for example, a backup storage system, as some storage system that makes use of the synchronously replicated dataset, and so on. In fact, synchronous replication can be leveraged to distribute copies of a dataset closer to some rack of servers, for better local read performance. One such case is smaller top-of-rack storage systems symmetrically replicated to larger storage systems that are centrally located in the data center or campus and where those larger storage systems are more carefully managed for reliability or are connected to external networks for asynchronous replication or backup services.

In the example depicted in FIG. 4, the depicted storage systems (402, 404, 406) may also identify a portion of the dataset that is not being asynchronously replicated to the target storage system by any of the other storage systems and asynchronously replicate, to the target storage system, the portion of the dataset that is not being asynchronously replicated to the target storage system by any of the other storage systems, wherein the two or more storage systems collectively replicate the entire dataset to the target storage system. In such a way, the work associated with asynchronously replicating a particular dataset may be split amongst the members of a pod, such that each storage system in a pod is only responsible for asynchronously replicating a subset of a dataset to the target storage system.

In the example depicted in FIG. 4, the depicted storage systems (402, 404, 406) may also detach from the pod, such that the particular storage system that detaches from the pod is no longer included in the set of storage systems across which the dataset is synchronously replicated. For example, if storage system (404) in FIG. 4 detached from the pod (430) illustrated in FIG. 4, the pod (430) would only include storage systems (402, 406) as the storage systems across which the dataset (426) that is included in the pod (430) would be synchronously replicated across. In such an example, detaching the storage system from the pod could also include removing the dataset from the particular storage system that detached from the pod. Continuing with the example where the storage system (404) in FIG. 4 detached from the pod (430) illustrated in FIG. 4, the dataset (426) that is included in the pod (430) could be deleted or otherwise removed from the storage system (404).

Readers will appreciate that there are a number of unique administrative capabilities enabled by the pod model that can further be supported. Also, the pod model itself introduces some issues that can be addressed by an implementation. For example, when a storage system is offline for a

pod, but is otherwise running, such as because an interconnect failed and another storage system for the pod won out in mediation, there may still be a desire or need to access the offline pod's dataset on the offline storage system. One solution may be simply to enable the pod in some detached mode and allow the dataset to be accessed. However, that solution can be dangerous and that solution can cause the pod's metadata and data to be much more difficult to reconcile when the storage systems do regain communication. Furthermore, there could still be a separate path for hosts to access the offline storage system as well as the still online storage systems. In that case, a host might issue I/O to both storage systems even though they are no longer being kept in sync, because the host sees target ports reporting volumes with the same identifiers and the host I/O drivers presume it sees additional paths to the same volume. This can result in fairly damaging data corruption as reads and writes issued to both storage systems are no longer consistent even though the host presumes they are. As a variant of this case, in a clustered application, such as a shared storage clustered database, the clustered application running on one host might be reading or writing to one storage system and the same clustered application running on another host might be reading or writing to the "detached" storage system, yet the two instances of the clustered application are communicating between each other on the presumption that the dataset they each see is entirely consistent for completed writes. Since they aren't consistent, that presumption is violated and the application's dataset (e.g., the database) can quickly end up being corrupted.

One way to solve both of these problems is to allow for an offline pod, or perhaps a snapshot of an offline pod, to be copied to a new pod with new volumes that have sufficiently new identities that host I/O drivers and clustered applications won't confuse the copied volumes as being the same as the still online volumes on another storage system. Since each pod maintains a complete copy of the dataset, which is crash consistent but perhaps slightly different from the copy of the pod dataset on another storage system, and since each pod has an independent copy of all data and metadata needed to operate on the pod content, it is a straightforward problem to make a virtual copy of some or all volumes or snapshots in the pod to new volumes in a new pod. In a logical extent graph implementation, for example, all that is needed is to define new volumes in a new pod which reference logical extent graphs from the copied pod associated with the pod's volumes or snapshots, and with the logical extent graphs being marked as copy on write. The new volumes should be treated as new volumes, similarly to how volume snapshots copied to a new volume might be implemented. Volumes may have the same administrative name, though within a new pod namespace. But, they should have different underlying identifiers, and differing logical unit identifiers from the original volumes.

In some cases it may be possible to use virtual network isolation techniques (for example, by creating a virtual LAN in the case of IP networks or a virtual SAN in the case of fiber channel networks) in such a way that isolation of volumes presented to some interfaces can be assured to be inaccessible from host network interfaces or host SCSI initiator ports that might also see the original volumes. In such cases, it may be safe to provide the copies of volumes with the same SCSI or other storage identifiers as the original volumes. This could be used, for example, in cases where the applications expect to see a particular set of storage identifiers in order to function without an undue burden in reconfiguration.

Some of the techniques described herein could also be used outside of an active fault context to test readiness for handling faults. Readiness testing (sometimes referred to as “fire drills”) is commonly required for disaster recovery configurations, where frequent and repeated testing is considered a necessity to ensure that most or all aspects of a disaster recovery plan are correct and account for any recent changes to applications, datasets, or changes in equipment. Readiness testing should be non-disruptive to current production operations, including replication. In many cases the real operations can’t actually be invoked on the active configuration, but a good way to get close is to use storage operations to make copies of production datasets, and then perhaps couple that with the use of virtual networking, to create an isolated environment containing all data that is believed necessary for the important applications that must be brought up successfully in cases of disasters. Making such a copy of a synchronously replicated (or even an asynchronously replicated) dataset available within a site (or collection of sites) that is expected to perform a disaster recovery readiness test procedure and then starting the important applications on that dataset to ensure that it can startup and function is a great tool, since it helps ensure that no important parts of the application datasets were left out in the disaster recovery plan. If necessary, and practical, this could be coupled with virtual isolated networks coupled perhaps with isolated collection of physical or virtual machines, to get as close as possible to a real world disaster recovery takeover scenario. Virtually copying a pod (or set of pods) to another pod as a point-in-time image of the pod datasets immediately creates an isolated dataset that contains all the copied elements and that can then be operated on essentially identically to the originally pods, as well as allowing isolation to a single site (or a few sites) separately from the original pod. Further, these are fast operations and they can be torn down and repeated easily allowing testing to repeated as often as is desired.

Some enhancements could be made to get further toward perfect disaster recovery testing. For example, in conjunction with isolated networks, SCSI logical unit identities or other types of identities could be copied into the target pod so that the test servers, virtual machines, and applications see the same identities. Further, the administrative environment of the servers could be configured to respond to requests from a particular virtual set of virtual networks to respond to requests and operations on the original pod name so scripts don’t require use of test-variants with alternate “test” versions of object names. A further enhancement can be used in cases where the host-side server infrastructure that will take over in the case of a disaster takeover can be used during a test. This includes cases where a disaster recovery data center is completely stocked with alternative server infrastructure that won’t generally be used until directed to do so by a disaster. It also includes cases where that infrastructure might be used for non-critical operations (for example, running analytics on production data, or simply supporting application development or other functions which may be important but can be halted if needed for more critical functions). Specifically, host definitions and configurations and the server infrastructure that will use them can be set up as they will be for an actual disaster recovery takeover event and tested as part of disaster recovery takeover testing, with the tested volumes being connected to these host definitions from the virtual pod copy used to provide a snapshot of the dataset. From the standpoint of the storage systems involved, then, these host definitions and configurations used for testing, and the

volume-to-host connection configurations used during testing, can be reused when an actual disaster takeover event is triggered, greatly minimizing the configuration differences between the test configuration and the real configuration that will be used in case of a disaster recovery takeover.

In some cases it may make sense to move volumes out of a first pod and into a new second pod including just those volumes. The pod membership and high availability and recovery characteristics can then be adjusted separately, and administration of the two resulting pod datasets can then be isolated from each other. An operation that can be done in one direction should also be possible in the other direction. At some point, it may make sense to take two pods and merge them into one so that the volumes in each of the original two pods will now track each other for storage system membership and high availability and recovery characteristics and events. Both operations can be accomplished safely and with reasonably minimal or no disruption to running applications by relying on the characteristics suggested for changing mediation or quorum properties for a pod which were discussed in an earlier section. With mediation, for example, a mediator for a pod can be changed using a sequence consisting of a step where each storage system in a pod is changed to depend on both a first mediator and a second mediator and each is then changed to depend only on the second mediator. If a fault occurs in the middle of the sequence, some storage systems may depend on both the first mediator and the second mediator, but in no case will recovery and fault handling result in some storage systems depending only on the first mediator and other storage systems only depending on the second mediator. Quorum can be handled similarly by temporarily depending on winning against both a first quorum model and a second quorum model in order to proceed to recovery. This may result in a very short time period where availability of the pod in the face of faults depend on additional resources, thus reducing potential availability, but this time period is very short and the reduction in availability is often very little. With mediation, if the change in mediator parameters is nothing more than the change in the key used for mediation and the mediation service used is the same, then the potential reduction in availability is even less, since it now depends only on two calls to the same service versus one call to that service, and rather than separate calls to two separate services.

Readers will note that changing the quorum model may be quite complex. An additional step may be necessary where storage systems will participate in the second quorum model but won’t depend on winning in that second quorum model, which is then followed by the step of also depending on the second quorum model. This may be necessary to account for the fact that if only one system has processed the change to depend on the quorum model, then it will never win quorum since there will never be a majority. With this model in place for changing the high availability parameters (mediation relationship, quorum model, takeover preferences), we can create a safe procedure for these operations to split a pod into two or to join two pods into one. This may require adding one other capability: linking a second pod to a first pod for high availability such that if two pods include compatible high availability parameters the second pod linked to the first pod can depend on the first pod for determining and instigating detach-related processing and operations, offline and in-sync states, and recovery and resynchronization actions.

To split a pod into two, which is an operation to move some volumes into a newly created pod, a distributed

operation may be formed that can be described as: form a second pod into which we will move a set of volumes which were previously in a first pod, copy the high availability parameters from the first pod into the second pod to ensure they are compatible for linking, and link the second pod to the first pod for high availability. This operation may be encoded as messages and should be implemented by each storage system in the pod in such a way that the storage system ensures that the operation happens completely on that storage system or does not happen at all if processing is interrupted by a fault. Once all in-sync storage systems for the two pods have processed this operation, the storage systems can then process a subsequent operation which changes the second pod so that it is no longer linked to the first pod. As with other changes to high availability characteristics for a pod, this involves first having each in-sync storage system change to rely on both the previous model (that model being that high availability is linked to the first pod) and the new model (that model being its own now independent high availability). In the case of mediation or quorum, this means that storage systems which processed this change will first depend on mediation or quorum being achieved as appropriate for the first pod and will additionally depend on a new separate mediation (for example, a new mediation key) or quorum being achieved for the second pod before the second pod can proceed following a fault that required mediation or testing for quorum. As with the previous description of changing quorum models, an intermediate step may set storage systems to participate in quorum for the second pod before the step where storage systems participate in and depend on quorum for the second pod. Once all in-sync storage systems have processed the change to depend on the new parameters for mediation or quorum for both the first pod and the second pod, the split is complete.

Joining a second pod into a first pod operates essentially in reverse. First, the second pod must be adjusted to be compatible with the first pod, by having an identical list of storage systems and by having a compatible high availability model. This may involve some set of steps such as those described elsewhere in this paper to add or remove storage systems or to change mediator and quorum models. Depending on implementation, it may be necessary only to reach an identical list of storage systems. Joining proceeds by processing an operation on each in-sync storage system to link the second pod to the first pod for high availability. Each storage system which processes that operation will then depend on the first pod for high availability and then the second pod for high availability. Once all in-sync storage systems for the second pod have processed that operation, the storage systems will then each process a subsequent operation to eliminate the link between the second pod and the first pod, migrate the volumes from the second pod into the first pod, and delete the second pod. Host or application dataset access can be preserved throughout these operations, as long as the implementation allows proper direction of host or application dataset modification or read operations to the volume by identity and as long as the identity is preserved as appropriate to the storage protocol or storage model (for example, as long as logical unit identifiers for volumes and use of target ports for accessing volumes are preserved in the case of SCSI).

Migrating a volume between pods may present issues. If the pods have an identical set of in-sync membership storage systems, then it may be straightforward: temporarily suspend operations on the volumes being migrated, switch control over operations on those volumes to controlling

software and structures for the new pod, and then resume operations. This allows for a seamless migration with continuous uptime for applications apart from the very brief operation suspension, provided network and ports migrate properly between pods. Depending on the implementation, suspending operations may not even be necessary, or may be so internal to the system that the suspension of operations has no impact. Copying volumes between pods with different in-sync membership sets is more of a problem. If the target pod for the copy has a subset of in-sync members from the source pod, this isn't much of a problem: a member storage system can be dropped safely enough without having to do more work. But, if the target pod adds in-sync member storage systems to the volume over the source pod, then the added storage systems must be synchronized to include the volume's content before they can be used. Until synchronized, this leaves the copied volumes distinctly different from the already synchronized volumes, in that fault handling differs and request handling from the not yet synced member storage systems either won't work or must be forwarded or won't be as fast because reads will have to traverse an interconnect. Also, the internal implementation will have to handle some volumes being in sync and ready for fault handling and others not being in sync.

There are other problems relating to reliability of the operation in the face of faults. Coordinating a migration of volumes between multi-storage-system pods is a distributed operation. If pods are the unit of fault handling and recovery, and if mediation or quorum or whatever means are used to avoid split-brain situations, then a switch in volumes from one pod with a particular set of state and configurations and relationships for fault handling, recovery, mediation and quorum to another then storage systems in a pod have to be careful about coordinating changes related to that handling for any volumes. Operations can't be atomically distributed between storage systems, but must be staged in some way. Mediation and quorum models essentially provide pods with the tools for implementing distributed transactional atomicity, but this may not extend to inter-pod operations without adding to the implementation.

Consider even a simple migration of a volume from a first pod to a second pod even for two pods that share the same first and second storage systems. At some point the storage systems will coordinate to define that the volume is now in the second pod and is no longer in the first pod. If there is no inherent mechanism for transactional atomicity across the storage systems for the two pods, then a naive implementation could leave the volume in the first pod on the first storage system and the second pod on the second storage system at the time of a network fault that results in fault handling to detach storage systems from the two pods. If pods separately determine which storage system succeeds in detaching the other, then the result could be that the same storage system detaches the other storage system for both pods, in which case the result of the volume migration recovery should be consistent, or it could result in a different storage system detaching the other for the two pods. If the first storage system detaches the second storage system for the first pod and the second storage system detaches the first storage system for the second pod, then recovery might result in the volume being recovered to the first pod on the first storage system and into the second pod on the second storage system, with the volume then running and exported to hosts and storage applications on both storage systems. If instead the second storage system detaches the first storage system for the first pod and first storage detaches the second storage system for the second pod, then recovery might

result in the volume being discarded from the second pod by the first storage system and the volume being discarded from the first pod by the second storage system, resulting in the volume disappearing entirely. If the pods a volume is being migrated between are on differing sets of storage systems, then things can get even more complicated.

A solution to these problems may be to use an intermediate pod along with the techniques described previously for splitting and joining pods. This intermediate pod may never be presented as visible managed objects associated with the storage systems. In this model, volumes to be moved from a first pod to a second pod are first split from the first pod into a new intermediate pod using the split operation described previously. The storage system members for the intermediate pod can then be adjusted to match the membership of storage systems by adding or removing storage systems from the pod, as necessary. Subsequently, the intermediate pod can be joined with the second pod.

For further explanation, FIG. 5 sets forth a flowchart illustrating an example method of applying a rate limit across a plurality of storage systems (504, 522) in accordance with some embodiments of the present disclosure. Although depicted in less detail, the storage systems depicted in FIG. 5 may be similar to the storage systems described above, as the storage systems (504, 522) may include any combination of the components contained in the storage systems described above. In the example method depicted in FIG. 5, the storage systems are (504, 522) considered to be paired storage systems (502) as some degree of coordination may be required between the storage systems. For example, a dataset may be synchronously replicated across the storage systems (504, 522) as described above with reference to FIG. 4.

The example method depicted in FIG. 5 includes determining (506) a rate limit for paired storage systems (504, 522). In the example depicted in FIG. 5, the rate limit specifies the cumulative amount of input/output ('I/O') operations that may be serviced by the paired storage systems (504, 522) during a predetermined period of time. The rate limit may be expressed, for example, in terms of an amount of data per unit of time (e.g., GB/s). Consider an example in which the rate limit for the paired storage systems (504, 522) is 10 GB/s. In such an example, all data written to either of the paired storage systems (504, 522) and all data read from either of the paired storage systems (504, 522) would count towards the rate limit. For example, if the first storage system (504) serviced 1 GB of write operations directed to the storage system (504) during a particular one second period, the first storage system (504) serviced 2.5 GB of read operations directed to the storage system (504) during the particular one second period, the second storage system (522) serviced 1.5 GB of write operations directed to the storage system (522) during a particular one second period, and the second storage system (522) serviced 2 GB of read operations directed to the storage system (522) during the particular one second period, the paired storage systems (504, 522) would have utilized 7 GB of its 10 GB limit during the particular one second period. Readers will appreciate that as the paired storage systems (504, 522) hit their rate limits, the servicing of I/O operations may be throttled (e.g., via queueing the I/O operations) to ensure that the rate limit is not exceeded.

Readers will appreciate the rate limit rate limit for paired storage systems (504, 522) may be distinct for one or more system-level entities. For example, each stretched volume (described above) may have a distinct rate limit for that stretched volume (e.g., the rate limit for a first stretched

volume may be 1 GB/s and the rate limit for a second stretched volume may be 1.5 GB/s). The rate limit for each system-level entity may be determined, for example, by dividing the total amount of I/O operations that may be serviced by the paired storage systems (504, 522) by the number of system-level entities that service I/O operations in a weighted or unweighted manner.

In the example method depicted in FIG. 5, determining (506) a rate limit for paired storage systems (504, 522) may be carried out, for example, by initially receiving the rate limit for paired storage systems (504, 522) as a configuration parameter and storing such a value within each of the storage systems (504, 522). In such an example, each storage system (504, 522) may periodically check its stored value to verify that the rate limit for the paired storage systems (504, 522) has not been changed.

The example method depicted in FIG. 5 also includes receiving (508), by a first storage system (504) from a second storage system (522), an amount (524) of I/O operations serviced by the second storage system (522) during a previous predetermined period of time. The amount (524) of I/O operations serviced by the second storage system (522) during a previous predetermined period of time may be expressed, for example, in terms of an amount of data that was read from or written to the second storage system (522) in order for the second storage system (522) to service I/O operations received by the second storage system (522) during a previous predetermined period of time. In the example method depicted in FIG. 5, the previous predetermined period of time may represent a polling interval that defines how frequently the second storage system (522) will send the first storage system (504) information describing the amount (524) of I/O operations serviced by the second storage system (522). Consider an example in which the second storage system (522) is configured to send the first storage system (504) information describing the amount (524) of I/O operations serviced by the second storage system (522) every 250 milliseconds. In such an example, every time that the second storage system (522) sends the first storage system (504) information describing the amount (524) of I/O operations serviced by the second storage system (522) during a previous predetermined period of time, the second storage system (522) will be sending the first storage system (504) information describing the amount (524) of I/O operations serviced by the second storage system (522) during the most recent 250 milliseconds. Stated differently, the previous predetermined period of time referenced above may be embodied as the most recent polling interval.

Readers will appreciate that the 'previous predetermined period of time' recited in step 508 and the 'predetermined period of time' recited in step 506 need not be the same duration. For example, if the rate limit for the paired storage systems (504, 522) is 10 GB/s and the first storage system (504) receives (508) information describing an amount (524) of I/O operations serviced by the second storage system (522) every 250 milliseconds, embodiments described herein may include making conversions (e.g., viewing the rate limit as 2.5 GB per 250 milliseconds) to create comparable values. In the example method depicted in FIG. 5, the first storage system (504) may receive (508) an amount (524) of I/O operations serviced by the second storage system (522) during a previous predetermined period of time from the second storage system (522) via one or more messages sent from the second storage system (522) to the first storage system (504) or in some other way.

The example method depicted in FIG. 5 also includes determining (510) whether the amount (524) of I/O operations serviced by the second storage system (522) during a previous predetermined period of time is less than half of the rate limit for the paired storage systems (504, 522). Continuing with the example described above where the rate limit for the paired storage systems (504, 522) is 10 GB/s (or 2.5 GB per 250 milliseconds), assume that the first storage system (504) received (508) information indicating that the second storage system (522) has serviced 1.0 GB during the most recent 250 millisecond period of time. In such an example, an affirmative (512) determination would be made that the amount (524) of I/O operations serviced by the second storage system (522) during a previous predetermined period of time is less than half of the rate limit for the paired storage systems (504, 522), thereby indicating that the second storage system (522) is not utilizing its fair share of the rate limit that would otherwise be reserved for use by the second storage system (522).

The example method depicted in FIG. 5 also includes, responsive to affirmatively (512) determining that the amount (524) of I/O operations serviced by the second storage system (522) during the previous predetermined period of time is less than half of the rate limit for the paired storage systems (504, 522), setting (516) a local rate limit for a next predetermined period of time for the first storage system (504) to the difference between the rate limit for the paired storage systems (504, 522) and the amount (524) of I/O operations serviced by the second storage system (522) during the previous predetermined period of time. The local rate limit for the first storage system (504) specifies the amount of I/O operations that may be serviced by the first storage system (504) during a predetermined period of time. The rate limit may be expressed, for example, in terms of an amount of data per unit of time (e.g., GB/s).

Continuing with the example described where the rate limit for the paired storage systems (504, 522) is 10 GB/s (or 2.5 GB per 250 milliseconds) and the first storage system (504) received (508) information indicating that the second storage system (522) has serviced 1.0 GB during the most recent 250 millisecond period of time, setting (516) a local rate limit for a next predetermined period of time for the first storage system (504) to the difference between the rate limit for the paired storage systems (504, 522) and the amount (524) of I/O operations serviced by the second storage system (522) during the previous predetermined period of time would result in the rate limit for the first storage system being set to a value of 1.5 GB for the next 250 millisecond interval. As such, because the second storage system (522) is not utilizing its fair share of the rate limit that would otherwise be reserved for use by the second storage system (522), the first storage system (504) may be given more than its fair share of the rate limit that would otherwise be reserved for use by the first storage system (504).

The example method depicted in FIG. 5 also includes, responsive to determining that the amount (524) of I/O operations serviced by the second storage system (522) during the previous predetermined period of time is not (514) less than half of the rate limit for the paired storage systems (504, 522), setting (518) a local rate limit for a next predetermined period of time for the first storage system (504) to half of the rate limit for the paired storage systems (504, 522). Continuing with the example described above where the rate limit for the paired storage systems (504, 522) is 10 GB/s (or 2.5 GB per 250 milliseconds), assume that the first storage system (504) received (508) information indicating that the second storage system (522) has serviced 1.5

GB during the most recent 250 millisecond period of time. In such an example, a determination would be made that the amount (524) of I/O operations serviced by the second storage system (522) during the previous predetermined period of time is not (514) less than half of the rate limit for the paired storage systems (504, 522), thereby indicating that the second storage system (522) is utilizing at least its fair share of the rate limit that would otherwise be reserved for use by the second storage system (522).

Continuing with the example described where the rate limit for the paired storage systems (504, 522) is 10 GB/s (or 2.5 GB per 250 milliseconds) and the first storage system (504) received (508) information indicating that the second storage system (522) has serviced 1.5 GB during the most recent 250 millisecond period of time, setting (518) a local rate limit for a next predetermined period of time for the first storage system (504) to half of the rate limit for the paired storage systems (504, 522) would result in the rate limit for the first storage system being set to a value of 1.25 GB for the next 250 millisecond interval. As such, because the second storage system (522) is utilizing at least its fair share of the rate limit that would otherwise be reserved for use by the second storage system (522), the first storage system (504) may not be given more than its fair share of the rate limit that would otherwise be reserved for use by the first storage system (504).

Although depicted in less detail, each storage system depicted in FIG. 4 may include a controller that services I/O operations directed to the storage system (504, 522). In such an example, each controller may include a plurality of nodes (e.g., a plurality of Non-uniform memory access (NUMA) nodes) that handle each handle some subset of the I/O operations serviced by each controller. In such an example, each node may participate in an exchange of information that mirrors the interactions of the storage systems described above. For example, a first node in a first storage system (504) may determine a local rate limit for the storage system (504) and may also receive, from a second node in the storage system (504), an amount of I/O operations serviced by the second node during a previous predetermined period of time. The first node may subsequently determine whether the amount of I/O operations serviced by the second node during the previous predetermined period of time is less than half of the local rate limit for the storage system (504). In response to determining that the amount of I/O operations serviced by the second node is less than half of the local rate limit for the storage system (504), the first node may set a local rate limit for a next predetermined period of time for the first node to the difference between the local rate limit for the storage system (504) and the amount of I/O operations serviced by the second node during the previous predetermined period of time. Alternatively; in response to determining that the amount of I/O operations serviced by the second node during the previous predetermined period of time is not less than half of the local rate limit for the storage system (504), the first node may set its local rate limit for a next predetermined period of time to half of the local rate limit of the storage system (504). The second node may similarly perform similar steps.

Readers will appreciate that through the use of the techniques described above, a hierarchy may essentially be created where a pair of storage systems negotiate their individual local rate limits based on the activity of the other storage system in the pair, and where pairs of nodes within an individual storage system similarly negotiate their local rate limits based on the activity of the other paired node and based on the local rate limit of the storage system that the

61

pair of nodes reside within. As such, the rate limit for a first node within the particular storage system may be set in dependence upon amount of I/O operations serviced by the a second node within the particular storage system during a previous predetermined period of time and the local rate limit for the particular storage system.

For further explanation, FIG. 6 sets forth a flowchart illustrating an additional example method of applying a rate limit across a plurality of storage systems (504, 522) in accordance with some embodiments of the present disclosure. The example method depicted in FIG. 6 is similar to the example method depicted in FIG. 5, as the example method depicted in FIG. 6 also includes determining (506) a rate limit for paired storage systems (504, 522), receiving (508) an amount (524) of I/O operations serviced by the second storage system (522) during a previous predetermined period of time, determining (510) whether the amount (524) of I/O operations serviced by the second storage system (522) during a previous predetermined period of time is less than half of the rate limit for the paired storage systems (504, 522), and setting (516, 518) a local rate limit for a next predetermined period of time for the first storage system (504).

The example method depicted in FIG. 6 also includes determining (602) a rate limit for paired storage systems (504, 522), the rate limit specifying the cumulative amount of I/O operations that may be serviced by the paired storage systems (504, 522) during a predetermined period of time. Determining (602) a rate limit for paired storage systems (504, 522) may be carried out, for example, by initially receiving the rate limit for paired storage systems (504, 522) as a configuration parameter and storing such a value within each of the storage systems (504, 522). In such an example, each storage system (504, 522) may periodically check its stored value to verify that the rate limit for the paired storage systems (504, 522) has not been changed.

The example method depicted in FIG. 6 also includes receiving (604), by the second storage system (522) from the first storage system (504), an amount (616) of I/O operations serviced by the first storage system (504) during a previous predetermined period of time. The amount (616) of I/O operations serviced by the first storage system (504) during a previous predetermined period of time may be expressed, for example, in terms of an amount of data that was read from or written to the first storage system (504) in order for the first storage system (504) to service I/O operations received by the first storage system (504) during a previous predetermined period of time. In the example method depicted in FIG. 6, the previous predetermined period of time may represent a polling interval that defines how frequently the first storage system (504) will send the second storage system (522) information describing the amount (616) of I/O operations serviced by the first storage system (504). Consider an example in which the first storage system (504) is configured to send the second storage system (522) information describing the amount (616) of I/O operations serviced by the first storage system (504) every 250 milliseconds. In such an example, every time that the first storage system (504) sends the second storage system (522) information describing the amount (616) of I/O operations serviced by the first storage system (504) during a previous predetermined period of time, the first storage system (504) will be sending the second storage system (522) information describing the amount (616) of I/O operations serviced by the first storage system (504) during the most recent 250

62

milliseconds. Stated differently, the previous predetermined period of time referenced above may be embodied as the most recent polling interval.

The example method depicted in FIG. 6 also includes determining (606) whether the amount of I/O operations serviced by the first storage system (504) during a previous predetermined period of time is less than half of the rate limit for the paired storage systems (504, 522). Continuing with the example described above where the rate limit for the paired storage systems (504, 522) is 10 GB/s (or 2.5 GB per 250 milliseconds), assume that the second storage system (522) received (604) information indicating that the first storage system (504) has serviced 1.0 GB during the most recent 250 millisecond period of time. In such an example, an affirmative (612) determination would be made that the amount (616) of I/O operations serviced by the first storage system (504) during a previous predetermined period of time is less than half of the rate limit for the paired storage systems (504, 522), thereby indicating that the first storage system (504) is not utilizing its fair share of the rate limit that would otherwise be reserved for use by the first storage system (504).

The example method depicted in FIG. 6 also includes, responsive to affirmatively (612) determining that the amount (616) of I/O operations serviced by the first storage system (504) during the previous predetermined period of time is less than half of the rate limit for the paired storage systems (504, 522), setting (608) a local a rate limit for a next predetermined period of time for the second storage system (522) to the difference between the rate limit for the paired storage systems (504, 522) and the amount (616) of I/O operations serviced by the first storage system (504) during the previous predetermined period of time. The local rate limit for the second storage system (522) specifies the amount of I/O operations that may be serviced by the second storage system (522) during a predetermined period of time. The rate limit may be expressed, for example, in terms of an amount of data per unit of time (e.g., GB/s).

Continuing with the example described where the rate limit for the paired storage systems (504, 522) is 10 GB/s (or 2.5 GB per 250 milliseconds) and the second storage system (522) received (604) information indicating that the first storage system (504) has serviced 1.0 GB during the most recent 250 millisecond period of time, setting (608) a local rate limit for a next predetermined period of time for the second storage system (522) to the difference between the rate limit for the paired storage systems (504, 522) and the amount (616) of I/O operations serviced by the first storage system (504) during the previous predetermined period of time would result in the rate limit for the second storage system (522) being set to a value of 1.5 GB for the next 250 millisecond interval. As such, because the first storage system (504) is not utilizing its fair share of the rate limit that would otherwise be reserved for use by the first storage system (504), the second storage system (522) may be given more than its fair share of the rate limit that would otherwise be reserved for use by the second storage system (522).

The example method depicted in FIG. 6 also includes, responsive to determining that the amount (616) of I/O operations serviced by the first storage system (504) during the previous predetermined period of time is not (614) less than half of the rate limit for the paired storage systems (504, 522), setting (610) a local rate limit for a next predetermined period of time for the second storage system (522) to half of the rate limit for the paired storage systems (504, 522). Continuing with the example described above where the rate limit for the paired storage systems (504, 522) is 10 GB/s (or

2.5 GB per 250 milliseconds), assume that the second storage system (522) received (604) information indicating that the first storage system (504) has serviced 1.5 GB during the most recent 250 millisecond period of time. In such an example, a determination would be made that the amount (616) of I/O operations serviced by the first storage system (504) during the previous predetermined period of time is not (614) less than half of the rate limit for the paired storage systems (504, 522), thereby indicating that the first storage system (504) is utilizing at least its fair share of the rate limit that would otherwise be reserved for use by the first storage system (504).

Continuing with the example described where the rate limit for the paired storage systems (504, 522) is 10 GB/s (or 2.5 GB per 250 milliseconds) and the second storage system (522) received (604) information indicating that the first storage system (504) has serviced 1.5 GB during the most recent 250 millisecond period of time, setting (610) a local rate limit for a next predetermined period of time for the second storage system (522) to half of the rate limit for the paired storage systems (504, 522) would result in the rate limit for the second storage system (522) being set to a value of 1.25 GB for the next 250 millisecond interval. As such, because the first storage system (504) is utilizing at least its fair share of the rate limit that would otherwise be reserved for use by the first storage system (504), the second storage system (522) may not be given more than its fair share of the rate limit that would otherwise be reserved for use by the second storage system (522).

For further explanation, FIG. 7 sets forth a flowchart illustrating an additional example method of applying a rate limit across a plurality of storage systems (504, 522) in accordance with some embodiments of the present disclosure. The example method depicted in FIG. 7 is similar to the example methods depicted in FIG. 5 and FIG. 6, as the example method depicted in FIG. 6 also includes determining (506) a rate limit for paired storage systems (504, 522), receiving (508) an amount (524) of I/O operations serviced by the second storage system (522) during a previous predetermined period of time, determining (510) whether the amount (524) of I/O operations serviced by the second storage system (522) during a previous predetermined period of time is less than half of the rate limit for the paired storage systems (504, 522), setting (516, 518) a local rate limit for a next predetermined period of time for the first storage system (504), receiving (604) an amount (616) of I/O operations serviced by the first storage system (504) during a previous predetermined period of time, determining (606) whether the amount of I/O operations serviced by the first storage system (504) during a previous predetermined period of time is less than half of the rate limit for the paired storage systems (504, 522), and setting (608, 610) a local rate limit for a next predetermined period of time for the second storage system (522).

The example method depicted in FIG. 7 also includes sending (702), from the first storage system (504) to the second storage system (522) at a predetermined interval, information describing the amount (524) of I/O operations serviced by the first storage system (504) during a previous predetermined period of time. The predetermined interval may be a configurable parameter and may represent the frequency at which each of the paired storage systems (504, 522) should send each other information describing the amount of I/O operations that a particular storage system serviced during a previous predetermined period of time.

Consider the example described above where the predetermined interval is 250 milliseconds. In such an example,

every 250 milliseconds the first storage system (504) may send (702), to the second storage system (522), information describing the amount (524) of I/O operations serviced by the first storage system (504) during the previous 250 milliseconds. The first storage system (504) may send (702) such information to the second storage system (522), for example, via one or more messages, by writing (e.g., via RDMA) such information into a memory location monitored by the second storage system (522), or in other ways. In order to facilitate the regular exchange of such information between the paired storage systems (504, 522), the example method depicted in FIG. 7 similarly includes sending (704), from the second storage system (522) to the first storage system (504) at the predetermined interval, information describing an amount (616) of I/O operations serviced by the second storage system (522) during the previous predetermined period of time.

Readers will appreciate that some embodiments of the present disclosure can include modifying the predetermined interval. In such an example, a determination may be made that the predetermined interval is too long or too short, and as such, the predetermined interval may need to be modified. As such, the predetermined interval represents a configurable parameter that may be tweaked, as necessary.

Readers will further appreciate that in some embodiments the sum of the local rate limit for the first storage system (504) and the local rate limit for the second storage system (522) can exceed the rate limit specifying the cumulative amount of I/O operations that may be serviced by the paired storage systems (504, 522) during the predetermined period of time. Consider the example described above where the rate limit for the paired storage systems (504, 522) is 10 GB/s. In such an example, assume that the predetermined interval upon which the storage systems (504, 522) exchange information describing the amount of I/O operations serviced by each storage system (504, 522) during a previous predetermined interval is set to a value of one second. Further assume that neither storage system (504, 522) has processed any I/O operations during the most recent 1 second interval. In such an example, as described in greater detail, each storage system will set its local rate limit to 10 GB for the next 1 second interval, meaning that the storage systems (504, 522) could collectively service 20 GB of I/O operations in the next second, which is well above the 10 GB/s rate limit for the paired storage systems (504, 522). Readers will appreciate that in such an example, however, the sum of the local rate limit for the first storage system (504) and the local rate limit for the second storage system (522) will quickly be adjusted to comply with the rate limit specifying the cumulative amount of I/O operations that may be serviced by the paired storage systems (504, 522) after a subsequent exchange of information between the paired storage systems (504, 522).

For further explanation, FIG. 8 sets forth a flowchart illustrating an additional example method of applying a rate limit across a plurality of storage systems (504, 522) in accordance with some embodiments of the present disclosure. The example method depicted in FIG. 6 is similar to the example method depicted in FIGS. 5-7, as the example method depicted in FIG. 8 also includes determining (506) a rate limit for paired storage systems (504, 522), receiving (508) an amount (524) of I/O operations serviced by the second storage system (522) during a previous predetermined period of time, determining (510) whether the amount (524) of I/O operations serviced by the second storage system (522) during a previous predetermined period of time is less than half of the rate limit for the paired storage

systems (504, 522), and setting (516, 518) a local rate limit for a next predetermined period of time for the first storage system (504). Readers will further appreciate that the steps introduced in FIG. 8 could also be included in embodiments where the exchange of information is bi-directional as explicitly depicted in FIG. 6 and FIG. 7.

Readers will appreciate that the steps described above may be performed iteratively in accordance with a predetermined interval, such that the paired storage systems (504, 522) are repeatedly exchanging information describing the amount of I/O operations that they have processed and the paired storage systems (504, 522) are repeatedly adjusting their local rate limits. At some point in time, however, one of the paired storage systems (504, 522) may crash or otherwise become unreachable. As such, the example method depicted in FIG. 8 also includes detecting (802) that the first storage system (504) has not received information describing the amount (524) of I/O operations serviced by the second storage system (522) during the previous predetermined period of time. Detecting (802) that the first storage system (504) has not received information describing the amount (524) of I/O operations serviced by the second storage system (522) during the previous predetermined period of time may be carried out, for example, by determining that no messages have been received from the second storage system (522) within an expected period of time (e.g., the predetermined interval). In response to detecting (802) that the first storage system (504) has not received information describing the amount (524) of I/O operations serviced by the second storage system (522) during the previous predetermined period of time, the first storage system (504) may assume that the second storage system (522) is offline or otherwise unreachable. As such, the example method depicted in FIG. 8 also includes setting (804) the local rate limit for the next predetermined period of time for the first storage system (504) to the rate limit for the paired storage systems (504, 522).

Readers will appreciate that although the embodiments described above relate to embodiments where two storage systems are paired, in other embodiments additional storage systems may be paired. For example, some embodiments can include a system of paired storage systems, where each storage system is configured to determine a rate limit for the system of paired storage systems and receive, from each of the other storage systems in the system of paired storage systems, an amount of I/O operations serviced by the other storage system during a previous predetermined period of time in a manner as described above. Readers will appreciate that in some embodiments, rather than individually receiving an amount of I/O operations serviced by each other storage system during a previous predetermined period of time, each storage system may simply receive the cumulative amount of I/O operations serviced by each of the other storage systems in the system during a previous predetermined period of time. Each storage system, in the embodiments described herein, may also be configured to determine whether the cumulative amount of I/O operations serviced by each of the other storage systems in the system of paired storage systems during the previous predetermined period of time is less than a fair share of I/O bandwidth for the other storage systems in the system of paired storage systems.

In such embodiments, the 'fair share' of I/O bandwidth for each individual storage system would be equal to the rate limit for the system of paired storage systems divided by the number of storage systems in the system of paired storage systems. For example, if the system of paired storage systems included 10 storage systems and the rate limit for

the system of paired storage systems was 10 GB/s, the 'fair share' of I/O bandwidth for each individual storage system would be 1 GB/s. As such, the fair share of I/O bandwidth for the other storage systems in the system of paired storage systems would be equal to the sum of the fair share of I/O bandwidth for each other storage system in the system of paired storage systems (9 GB/s in the example described in the preceding sentences).

In such embodiments, each storage system may be configured to set its local rate limit for a next predetermined period of time to the difference between the rate limit for the system of paired storage systems and the cumulative amount of I/O operations serviced by each of the other storage systems in the system of paired storage systems during the previous predetermined period of time. Each storage system may set its local rate limit in such a way in response to affirmatively determining that the cumulative amount of I/O operations serviced by each of the other storage systems in the system of paired storage systems during the previous predetermined period of time is less than the fair share of I/O bandwidth for the other storage systems in the system of paired storage systems. If a storage system determines that the cumulative amount of I/O operations serviced by each of the other storage systems in the system of paired storage systems during the previous predetermined period of time is not less than the fair share of I/O bandwidth for the other storage systems in the system of paired storage systems, however, each storage system may be configured to set its local rate limit for a next predetermined period of time to a fair share of I/O bandwidth for each storage system.

Readers will appreciate that while the examples described in the preceding paragraphs relates to embodiments where a 'fair share' of I/O bandwidth for each storage system is utilized, alternative embodiments may not include the notion of a 'fair share' but instead may rely on allowing for different storage systems to have access to unequal portions of the total rate limit for the system of storage systems. In such examples, each storage system may be configured to determine whether the cumulative amount of I/O operations serviced by each of the other storage systems in the system during the previous predetermined period of time is less than a 'configurable share' of I/O bandwidth for the other storage systems in the system of paired storage systems. In such example, each of the storage systems may be associated with a predetermined share of I/O bandwidth that is essentially reserved for the storage system, where the predetermined share of I/O bandwidth is different for at least two of the storage systems. In such an example, the 'configurable share' of I/O bandwidth for the other storage systems in the system of paired storage systems may be equal to the sum of the predetermined share of I/O bandwidth for each other storage system. In such an example, each storage system may be configured to set its local rate limit for a next predetermined period of time to the difference between the rate limit for the system of paired storage systems and the cumulative amount of I/O operations serviced by each of the other storage systems in the system during the previous predetermined period of time in response to affirmatively determining that the cumulative amount of I/O operations serviced by each of the other storage systems in the system during the previous predetermined period of time is less than a configurable share of I/O bandwidth for the other storage systems in the system of paired storage systems. Likewise, each storage system may be configured to set its local rate limit for a next predetermined period of time to its predetermined share of I/O bandwidth in response to determining that the cumulative amount of I/O operations serviced by

each of the other storage systems in the system during the previous predetermined period of time is not less than the configurable share of I/O bandwidth for the other storage systems in the system of paired storage systems. Readers will appreciate that in other embodiments, some hybrid between the 'fair share' and 'configurable share' embodiments may be implemented.

In the embodiments described above, each storage system may be configured to perform many of the same steps as described above such as, for example, sending information describing the amount of I/O operations serviced by the storage system during a previous predetermined period of time to at least one of the other storage systems at a predetermined interval. Likewise, each storage system may be configured to modify the predetermined interval as described above. Each storage system may similarly be configured to detecting that the storage system has not received information describing the amount of I/O operations serviced by at least one other storage system during the previous predetermined period of time and may assume (for the purposes of making the determinations described above) that the at least one other storage system did not service any I/O operations during the previous predetermined period of time.

Readers will appreciate that although the examples described above are depicted in a way where steps occur in a certain order, no such ordering of steps is required unless explicitly stated. Example embodiments are described largely in the context of a fully functional computer system. Readers of skill in the art will recognize, however, that the present disclosure also may be embodied in a computer program product disposed upon computer readable storage media for use with any suitable data processing system. Such computer readable storage media may be any storage medium for machine-readable information, including magnetic media, optical media, or other suitable media. Examples of such media include magnetic disks in hard drives or diskettes, compact disks for optical drives, magnetic tape, and others as will occur to those of skill in the art. Persons skilled in the art will immediately recognize that any computer system having suitable programming means will be capable of executing the steps of the method as embodied in a computer program product. Persons skilled in the art will recognize also that, although some of the example embodiments described in this specification are oriented to software installed and executing on computer hardware, nevertheless, alternative embodiments implemented as firmware or as hardware are well within the scope of the present disclosure.

Embodiments can include be a system, a method, and/or a computer program product. The computer program product may include a computer readable storage medium (or media) having computer readable program instructions thereon for causing a processor to carry out aspects of the present disclosure.

The computer readable storage medium can be a tangible device that can retain and store instructions for use by an instruction execution device. The computer readable storage medium may be, for example, but is not limited to, an electronic storage device, a magnetic storage device, an optical storage device, an electromagnetic storage device, a semiconductor storage device, or any suitable combination of the foregoing. A non-exhaustive list of more specific examples of the computer readable storage medium includes the following: a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory

(EPROM or Flash memory), a static random access memory (SRAM), a portable compact disc read-only memory (CD-ROM), a digital versatile disk (DVD), a memory stick, a floppy disk, a mechanically encoded device such as punch-cards or raised structures in a groove having instructions recorded thereon, and any suitable combination of the foregoing. A computer readable storage medium, as used herein, is not to be construed as being transitory signals per se, such as radio waves or other freely propagating electromagnetic waves, electromagnetic waves propagating through a waveguide or other transmission media (e.g., light pulses passing through a fiber-optic cable), or electrical signals transmitted through a wire.

Computer readable program instructions described herein can be downloaded to respective computing/processing devices from a computer readable storage medium or to an external computer or external storage device via a network, for example, the Internet, a local area network, a wide area network and/or a wireless network. The network may comprise copper transmission cables, optical transmission fibers, wireless transmission, routers, firewalls, switches, gateway computers and/or edge servers. A network adapter card or network interface in each computing/processing device receives computer readable program instructions from the network and forwards the computer readable program instructions for storage in a computer readable storage medium within the respective computing/processing device.

Computer readable program instructions for carrying out operations of the present disclosure may be assembler instructions, instruction-set-architecture (ISA) instructions, machine instructions, machine dependent instructions, microcode, firmware instructions, state-setting data, or either source code or object code written in any combination of one or more programming languages, including an object oriented programming language such as Smalltalk, C++ or the like, and conventional procedural programming languages, such as the "C" programming language or similar programming languages. The computer readable program instructions may execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user's computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider). In some embodiments, electronic circuitry including, for example, programmable logic circuitry, field-programmable gate arrays (FPGA), or programmable logic arrays (PLA) may execute the computer readable program instructions by utilizing state information of the computer readable program instructions to personalize the electronic circuitry, in order to perform aspects of the present disclosure.

Aspects of the present disclosure are described herein with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems), and computer program products according to some embodiments of the disclosure. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer readable program instructions.

These computer readable program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data pro-

cessing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks. These 5 computer readable program instructions may also be stored in a computer readable storage medium that can direct a computer, a programmable data processing apparatus, and/or other devices to function in a particular manner, such that the computer readable storage medium having instructions 10 stored therein comprises an article of manufacture including instructions which implement aspects of the function/act specified in the flowchart and/or block diagram block or blocks.

The computer readable program instructions may also be 15 loaded onto a computer, other programmable data processing apparatus, or other device to cause a series of operational steps to be performed on the computer, other programmable apparatus or other device to produce a computer implemented process, such that the instructions which execute on 20 the computer, other programmable apparatus, or other device implement the functions/acts specified in the flowchart and/or block diagram block or blocks.

The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible 25 implementations of systems, methods, and computer program products according to various embodiments of the present disclosure. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of instructions, which comprises one or more 30 executable instructions for implementing the specified logical function(s). In some alternative implementations, the functions noted in the block may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the 35 blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by 40 special purpose hardware-based systems that perform the specified functions or acts or carry out combinations of special purpose hardware and computer instructions.

What is claimed is:

1. A method comprising:

determining a rate limit for paired storage systems, the rate limit specifying the cumulative amount of input/output ('I/O') operations that may be serviced by the paired storage systems during a predetermined period of time;

determining whether the amount of I/O operations serviced by a second storage system during the previous predetermined period of time is less than a local rate limit for the second storage system;

responsive to determining that the amount of I/O operations 55 serviced by the second storage system during the previous predetermined period of time is less than the local rate limit for the second storage system, setting a local rate limit for a first storage system for a next predetermined period of time to the difference between 60 the rate limit for the paired storage systems and the amount of I/O operations serviced by the second storage system during the previous predetermined period of time.

2. The method of claim 1 further comprising:

determining whether the amount of I/O operations serviced by the first storage system during a previous

predetermined period of time is less than the local rate limit for the first storage system;

responsive to determining that the amount of I/O operations serviced by the first storage system during a previous predetermined period of time is less than the local rate limit for the first storage system, setting a local rate limit for a next predetermined period of time for the second storage system to the difference between the rate limit for the paired storage systems and the amount of I/O operations serviced by the first storage system during the previous predetermined period of time.

3. The method of claim 1 further comprising:

sending, from the first storage system to the second storage system at a predetermined interval, information describing the amount of I/O operations serviced by the first storage system during a previous predetermined period of time; and

sending, from the second storage system to the first storage system at the predetermined interval, information describing an amount of I/O operations serviced by the second storage system during the previous predetermined period of time.

4. The method of claim 2 wherein the sum of the local rate limit for the first storage system and the local rate limit for the second storage system exceeds the rate limit specifying the cumulative amount of I/O operations that may be serviced by the paired storage systems during the predetermined period of time.

5. The method of claim 1 further comprising:

detecting that the first storage system has not received information describing the amount of I/O operations serviced by the second storage system during the previous predetermined period of time; and

35 setting the local rate limit for the next predetermined period of time for the first storage system to the rate limit for the paired storage systems.

6. The method of claim 1 further comprising, within a particular storage system, setting the rate limit for a first node within the particular storage system in dependence upon amount of I/O operations serviced by the second node within the particular storage system during a previous predetermined period of time and the local rate limit for the particular storage system.

45 **7.** The method of claim 1 further comprising responsive to determining that the amount of I/O operations serviced by the first storage system during the previous predetermined period of time is not less than the local rate limit for the first storage system, setting a local rate limit for a next predetermined period of time for the second storage system to a predetermined portion of the rate limit for the paired storage systems.

8. An apparatus including a computer processor and a computer memory operatively coupled to the computer processor, the computer memory having disposed within it computer program instructions that, when executed by the computer processor, cause the apparatus to carry out the steps of:

determining a rate limit for paired storage systems, the rate limit specifying the cumulative amount of input/output ('I/O') operations that may be serviced by the paired storage systems during a predetermined period of time;

determining whether the amount of I/O operations serviced by a second storage system during the previous predetermined period of time is less than a local rate limit for the second storage system; and

responsive to determining that the amount of I/O operations serviced by the second storage system during the previous predetermined period of time is not less than the local rate limit for the second storage system, setting a local rate limit for a next predetermined period of time for the first storage system to a first predetermined portion of the rate limit for the paired storage systems.

9. The apparatus of claim **8** further comprising computer program instructions that, when executed by the computer processor, cause the apparatus to carry out the steps of:

determining whether the amount of I/O operations serviced by the first storage system during a previous predetermined period of time is less than a local rate limit for the first storage system; and

responsive to determining that the amount of I/O operations serviced by the first storage system during the previous predetermined period of time is not less than a local rate limit for the first storage system, setting a local rate limit for the second storage system for a next predetermined period of time to a second predetermined portion of the rate limit for the paired storage systems.

10. The apparatus of claim **8**, wherein the sum of the local rate limit for each storage system exceeds rate limit for the system of paired storage systems.

11. The apparatus of claim **8** further comprising computer program instructions that, when executed by the computer processor, cause the apparatus to carry out the step of:

detecting that information has not been received describing the amount of I/O operations serviced by a particular storage system during the previous predetermined period of time; and

assuming that the particular storage system did not service any I/O operations during the previous predetermined period of time.

12. The apparatus of claim **8** further comprising computer program instructions that, when executed by the computer processor, cause the apparatus to carry out the steps of, within a particular storage system, setting the rate limit for a first node within the particular storage system in dependence upon amount of I/O operations serviced by the a second node within the particular storage system during a previous predetermined period of time and the local rate limit for the particular storage system.

13. A computer program product for a system of paired storage systems, the system of paired storage systems including a plurality of storage systems, the computer program product disposed on a non-transitory computer readable medium that includes computer program instructions that, when executed by a computer processor, cause the computer processor to carry out the steps of:

determining a rate limit for the system of paired storage systems, the rate limit specifying the cumulative amount of input/output ('I/O') operations that may be serviced by the storage systems during a predetermined period of time;

receiving a cumulative amount of I/O operations serviced by each of the storage systems during a previous predetermined period of time;

determining, for a particular storage system, whether the cumulative amount of I/O operations serviced by each of the other storage systems in the system during the

previous predetermined period of time is less than a configurable share of I/O bandwidth for the other storage systems in the system of paired storage systems;

responsive to determining that the cumulative amount of I/O operations serviced by each of the other storage systems in the system during the previous predetermined period of time is less than a configurable share of I/O bandwidth for the other storage systems in the system of paired storage systems, setting a local rate limit for a next predetermined period of time for the particular storage system to the difference between the rate limit for the system of paired storage systems and the cumulative amount of I/O operations serviced by each of the other storage systems in the system during the previous predetermined period of time.

14. The computer program product of claim **13**, wherein each storage system sends, at a predetermined interval, information describing the amount of I/O operations serviced by the storage system during a previous predetermined period of time.

15. The computer program product of claim **13**, wherein the sum of the local rate limit for each storage system exceeds rate limit for the system of paired storage systems.

16. The computer program product of claim **13**, further comprising computer program instructions that, when executed by the computer processor, cause the computer processor to carry out the steps of:

detecting that the storage system has not received information describing the amount of I/O operations serviced by at least one other storage system during the previous predetermined period of time; and

assuming that the at least one other storage system did not service any I/O operations during the previous predetermined period of time.

17. The computer program product of claim **13**, wherein the predetermined share of I/O bandwidth for the storage system is different for at least two of the storage systems.

18. The computer program product of claim **17**, wherein the configurable share of I/O bandwidth for the other storage systems in the system of paired storage systems is equal to the sum of the predetermined share of I/O bandwidth for each other storage system.

19. The computer program product wherein each storage system further comprises computer program instructions that, when executed by the computer processor, cause each storage system to carry out the step of setting the rate limit for a first node within the particular storage system in dependence upon amount of I/O operations serviced by the a second node within the particular storage system during a previous predetermined period of time and the local rate limit for the particular storage system.

20. The computer program product of claim **13** further comprising responsive to determining that the cumulative amount of I/O operations serviced by each of the other storage systems in the system during the previous predetermined period of time of time is not less than the configurable share of I/O bandwidth for the other storage systems in the system of paired storage systems, setting a local rate limit for a next predetermined period of time to a predetermined share of I/O bandwidth for the particular storage system.